Published in final edited form as:

J Comb Optim. 2008 April 1; 15(3): 242–256. doi:10.1007/s10878-007-9101-5.

Tracing ‘driver’ versus ‘modulator’ information flow throughout large-scale, task-related neural circuitry

Linda Hermer-Vazquez

University of Florida, UF Box 112250, Gainesville, FL 32611, USA

Linda Hermer-Vazquez: lindahv@ufl.edu

Abstract

Primary objective—To determine the relative uses of neural action potential (‘spike’) data versus local field potentials (LFPs) for modeling information flow through complex brain networks.

Hypothesis—The common use of LFP data, which are continuous and therefore more mathematically suited for spectral information-flow modeling techniques such as Granger causality analysis, can lead to spurious inferences about whether a given brain area ‘drives’ the spiking in a downstream area.

Experiment—We recorded spikes and LFPs from the forelimb motor cortex (M1) and the magnocellular red nucleus (mRN), which receives axon collaterals from M1 projection cells onto its distal dendrites, but not onto its perisomatic regions, as rats performed a skilled reaching task.

Results and implications—As predicted, Granger causality analysis on the LFPs—which are mainly composed of vector-summed dendritic currents—produced results that if conventionally interpreted would suggest that the M1 cells drove spike firing in the mRN, whereas analyses of spiking in the two recorded regions revealed no significant correlations. These results suggest that mathematical models of information flow should treat the sampled dendritic activity as more likely to reflect intrinsic dendritic and input-related processing in neural networks, whereas spikes are more likely to provide information about the output of neural network processing.

Keywords

Information flow; Directed coherence; Granger causality; Action potentials; Local field potentials

1 Introduction

A question of great interest in neuroscience is how information flows throughout the extraordinarily complex circuitry of the central nervous system to encode information or produce adaptive behavior. Studying this question has required researchers to simplify the problem considerably. One approach has involved the separation of efferent fibers into ‘drivers’ on the one hand, which carry the information to be relayed, and ‘modulators’ on the other hand, which more subtly alter how driver information is processed at the projection site (s) (Sherman and Guillery 1998). Originally, the concept was most clearly applied to sensory relays such as retinal ganglion cells (drivers) versus neuromodulatory systems such as basal forebrain cholinergic projections (modulators). More recently, however, the concept has been

expanded in ways making it more useful for the study of the inner circuitry of the CNS (Sherman and Guillery 2002; Reichova and Sherman 2004), e.g.:

- Driver cells project to the downstream cells somatically or perisomatically, whereas modulators often synapse onto distal dendrites.
- Drivers often elicit large, all-or-none EPSPs that rely on fast-acting, ionotropic (e.g. AMPA) receptors, whereas modulators elicit much smaller EPSPs and may depend substantially on slower-acting, metabotropic receptors.

The concept of drivers versus modulators is relevant to the choice of neural data to be used in modeling information flow throughout neural circuitry. Local field potentials usually and primarily consist of vector-summed dendritic activity (Mitzdorf 1985; Bullock 1997). As such, they can be thought of as providing insight into the input of a neural network, reflecting dendritic computations, intrinsic activity, and many modulatory processes. In contrast, action potentials ('spikes') are capable of reflecting the output of a neural network, because their myelin insulation allows them to carry information over long distances with little susceptibility to noise (Laughlin and Sejnowski 2003).

When studying information flow through neural networks, one question researchers often address concerns whether neurons in a given area drive the spiking of cells in a downstream area. An increasingly commonly used approach (e.g. Brovelli et al. 2004; Faselow et al. 2001) employs the concept of 'Granger causality': If knowing the past of the time series $x_i(t)$ improves the prediction of the time series $x_j(t)$, then $x_i(t)$ can be said to 'Granger-cause' $x_j(t)$ (Granger 1969). This concept was originally developed for the time domain and was later extended to the frequency domain (Geweke 1982). Also called 'directed coherence' (in contrast to the more familiar 'classical coherence' between two time series, which reveals stationary phase-locking but does not yield relative phase information), this autoregressive technique searches for correlations between the spectra of the two time series at specified lead and lag times. Although the method is Fourier-based, its stationarity can be partly overcome by incorporating a moving window into the analysis (making it an autoregressive, moving-average, or 'ARMA,' technique). As a spectral method, it is most easily applied to continuous or non-spiky discrete signals. For this reason, local field potentials or EEG signals rather than action potentials from two brain sites (or more, if 'conditional Granger causality' is being performed; Baccala and Sameshima 2001; Chen et al. 2006) are most often used for these analyses e.g. (Sharott et al. 2005; Gueguin et al. 2006; Kus et al. 2004).

There are many pitfalls to such analyses beyond the stationarity problem. The simplest type of error, and the one most commonly cited (e.g. Chen et al. 2006), is depicted in Fig. 1(a) and 1(b). If one is analyzing recordings from brain areas 1 and 2, and the analysis indicates that area 1 drives area 2 (Fig. 1(a)), information may actually be flowing *from area 1 to area 3 to area 2* (Fig. 1(b)), a different causal chain. A further type of error, which has seldom been noted and is the subject of this article, is that if one uses LFPs from areas 1 and 2 for the analyses, then *the actual information flow from area 1 to area 2—which is likely to be spikes—is not even used in the analysis*. This can lead to extreme errors of interpretation. For instance, as depicted in Fig. 1(c), if area 1 and area 3 are simultaneously activated, but LFPs only from areas 1 and 2 are analyzed, it may appear that the input-related processing in area 2's dendrites (sampled by the LFP wires implanted near there) both (a) came from area 1 and (b) is driving the spiking in area 2. On the contrary, projections from area 3, which may synapse near the cell bodies in area 2, may actually have been driving the spike activity of area 2. To take a concrete example, one means by which cortical areas project to one another is via outputs from layer 5 of one area to a higher-order thalamic nucleus to layer 5 of the target site (Sherman 2005). Suppose the target site is a multi-modal association area, and several of its cortical afferent sites are active, although not necessarily all of them are sending action potentials to

the multi-modal area. In this case, the experimenter could record the dendritic activity of one of the projecting sites and the dendritic activity of the multi-modal site, and make the incorrect inference that particular output site is driving the efferent, multimodal site.

This potential problem has led us to propose the following hypothesis, which incorporates the notion of LFPs and modulators on the one hand, and spikes and drivers on the other. If one performs Granger causality analysis on LFPs from an area (area 1) that projects onto the *distal dendrites* of neurons in area 2, area 1 likely *modulates* but does not drive spiking in area 2. However, a strictly LFP-based analysis would spuriously indicate that area 1 ‘drove’ area 2. In contrast, a third area (area 3 in the diagram) that synapses somatically or perisomatically onto area 2’s neurons may play a much larger role in driving the spike activity in that area (area 2).

Testing this hypothesis requires considerable anatomical knowledge of the areas involved in the study. We implanted rats with multi-electrodes in their forelimb sensorimotor cortex (caudal M1) and their magnocellular red nucleus (mRN), and recorded unit activity (spikes) as well as local field potentials (LFPs) in both areas simultaneously. Our choice of these recording sites was based on the following anatomical logic (see Fig. 2). In rats, M1 pyramidal cells project to the spinal cord as part of the corticospinal tract, but they send axon collaterals to the *distal dendrites* of mRN cells (Keifer and Lustig 2000). In contrast, cells of the interpositus nucleus of the cerebellum (one of the deep cerebellar nuclei that play a major role in projecting cerebellar output to distal sites) synapse directly onto the somas of mRN neurons, as well as onto their proximal and distal dendrites (Keifer and Lustig 2000); this has also been found in the primate red nucleus (Ralston 1994b). Correspondingly, loss of the normal tonic and phasic cerebellar cortical inhibition on interpositus nucleus projection cells leads to a direct increase in mRN excitability and spike firing (Tarnecki 2003; Tarnecki et al. 2001; Thompson et al. 1997), demonstrating the substantial role that interpositus inputs play in determining mRN spiking. In general, synapses onto a cell’s soma or perisomatic region exert more control over that cell’s spike activity, whereas dendritic input, particularly if at the distal dendrites, plays a subtler, more modulatory role. Since unit recordings reflect spiking and LFPs primarily reflect dendritic activity, we predicted that our combination of information-flow analyses on spikes versus LFPs would demonstrate a *modulatory*, but *not a driving*, role for M1 in shaping mRN spike output. The experiments reported here are consistent with this prediction.

2 Methods

Behavioral training and testing

All experimental methods were approved by the University of Florida Institutional Animal Care and Use Committee. Four rats were taught a well studied reach-to-grasp-food task (Hermer-Vazquez et al. 2004; Whishaw and Pellis 1990; Kleim et al. 2002) prior to electrode implantation. In this task, rats learned to reach through a narrow plastic slot to grasp a 45 mg food pellet from one of two wells—the well contralateral to the reaching paw—located 13 mm beyond the slot on a plastic shelf at roughly the height of the rat’s snout (Fig. 3(a)). The slot’s width was such that the rat could place its nose directly behind the pellet in the well, thereby gaining information on the pellet’s odor quality and location, but still could not retrieve it with its tongue. Three fans at the rear of the chamber blew the pellet’s odor plume outside the reaching chamber so that rats had to sniff deeply to determine the pellet’s exact location, which varied on some trials. Once rats had asymptoted at a stable, high performance level, which took 5–8 days for rats in this cohort, the animals sniffed the target for 1–3 respiratory cycles before withdrawing their snout, lifting their reaching paw off the ground and maneuvering it through the slot to a position just beyond where their snout had been, allowing them to grasp the pellet and bring it to their mouth. The pellets were either banana- or chocolate-flavored and served as the reward when rats successfully retrieved them. Rats that reached an

asymptotic level then underwent electrode implantation surgery. After recovery, which took 1–2 weeks per rat, the animals were retrained in the task until their performance was similar to their pre-surgical level.

Videocoding

Using the video recordings made during the neural recording sessions (after rats had reached presurgical performance levels), we coded the times of the times of all task phases previously shown by our laboratory to correlate with significant spike rate modulation (Hermer-Vazquez et al. 2004). These phases are represented in a timeline in Fig. 3(b), and include the *Final Sniff* of the pellet before lifting the paw to reach, the moment the paw began to be lifted (*Lift*), the moment the rat made a complex upper-body and forearm motion in preparation to insert the paw through the slot (*Slot-align*), the moment the paw first touched the shelf with the pellet wells (*Onshelf*), and the time of *Contact* between the paw and the pellet. In addition, when one or two sniffs preceded the *Final Sniff*, we videocoded those times as well (*Sniff 1* and *Sniff 2*).

Electrode implantation

Rats were stereotaxically implanted with 2×8 electrode arrays in the caudal forelimb M1 that was contralateral to the reaching paw (0.5 mm anterior to bregma, 3.8 mm lateral, and approximately 1.6 mm deep) and 8-wire electrode bundles in the magnocellular red nucleus (5.9 mm posterior to bregma, 2.9 mm lateral angled medially at 16 degrees, and approximately 7.3 mm deep), also contralateral to the reaching paw. All electrodes were 50 μm at the tip, stainless steel, and Teflon-coated, fabricated by NB Labs (Denison, TX). The surgeries took place under a combination of ketamine/xylazine anesthesia, so that the rats were heavily anesthetized during the craniotomy portion of the surgery, and isoflurane gas, so that we could lighten the anesthesia precisely in order to see increased cellular activity as we descended, without causing the animals pain. At the end of the experiment all brains underwent histological analysis (Histology Tech Services, Gainesville, FL). All four rats in this study were found to have correctly placed M1 and mRN electrodes.

Neural recordings

Multi-single unit activity and local field potentials were recorded simultaneously using a Multi-neuron Acquisition Processor (MNAP) designed by Plexon, Inc. (Dallas, TX). Units were high-pass filtered above 6 kHz to remove 60 Hz artifact and sampled at 40 kHz. Waveform-based discrimination of single units was used, with the first two principal components analysis, K-means clustering and spike amplitude histograms employed in defining each cell. Isolation of single units was checked with autocorrelograms, and only units displaying the 1–2 ms notch in the center, corresponding to a single cell's refractory period, were used for analyses. Local field potentials were low-pass filtered at 100 Hz and sampled at 1 kHz.

Data analysis

All data analyses were performed with NeuroExplorer (Dallas, TX), Statistica (StatSoft, Tulsa, OK) and Matlab (The Mathworks, Natick, MA).

The first stage of the analyses was to construct perievent histograms for units from each recorded area, for each videocoded phase of the task. The purpose of this was to check for gross spike-firing relationships in cells across the two areas. For instance, if large populations of M1 and mRN cells fired above baseline for a particular phase, then this might be an early hint that one area was driving the other. To take another example, in our prior work (Hermer-Vazquez et al. 2004) we had found that many more cells in the mRN fired during the *Slot-align* phase than did M1 cells, leading to our hypothesis here that M1 cells would not be found

to drive or even necessarily correlate with mRN activity during that phase. These histograms used a 25 ms binwidth, with spikes/second on the y-axis and time relative to the videocoded event (e.g. *Lift*) on the x-axis (from 1 second before the event until 1 second afterward).

Next, we analysed LFPs from each area to check for directed coherence between them during any phase of the task. In previous work we had determined that mechanical artifacts such as those caused by the rat bumping its headcap against the testing chamber wall usually produced voltage fluctuations greater than ± 100 mV (Hermer-Vazquez and Hermer-Vazquez, unpublished data). In contrast, even relatively large-amplitude voltage fluctuations that were much more likely to be caused by neural activity alone, such as those seen as the quietly immobile rat falls asleep, seldom exceed this range using our equipment. Thus we excluded LFPs from further analysis if their voltage fluctuations exceeded this range. We removed DC offset through linear detrending, and we normalized the M1 and mRN time series so that the energy in the theta band was equal across the two areas on each trial. We analyzed LFP data from M1 and the mRN on a trial-by-trial basis, using a 100 ms analysis window that was moved across the trial's data in 10 ms increments. The trials were randomly selected except that trials with railing LFPs were omitted and randomly replaced. Although our sampling rate (1 kHz) would have permitted analysis of frequencies up to 500 Hz according to the Nyquist limit, our one-pole preamplifier filter began cutting off signals at ~ 100 Hz, so we restricted our analyses to the 0–100 Hz range.

Our method, described for ≥ 2 times series in Baccala and Sameshima (2001), involves writing the power spectral density for each of the two time series in such a way that they can be further analyzed for possible causality. If the past of $x_i(t)$ correlates with the present or future of $x_j(t)$, $x_i(t)$ may be driving $x_j(t)$. Briefly and essentially, in the spectral domain, the method involves computing the cross-spectra for the two time series and then quantitating, for a series of leads and lags (determined by the model order), whether the past of one series correlates with the present of another. Our normalization process led, for each trial, to two discrete time series (one for M1 and the other for the mRN) $x_i(n)$, $1 \leq i \leq N$, whose mutual relations can be described by a set of p (model order) matrices A_p of the autoregressive model (Fanselow et al. 2001). If $A(f) = I - A(f)$ for all $a_{ij}(f)$, knowing the past of $x_j(n)$ leads to no improvement of the prediction of $x_j(n)$'s present. Determining these relationships with an autoregressive, moving-average (ARMA) model as we did is desirable for two reasons. First, it reduces the amount of data needed for an accurate computation, and in our case we had to analyze the data from each trial, which was sometimes as short as 350 ms, separately. Second, the moving-window component partly deals with the stationarity problem. Note that driving by either of the two time series through feedforward or feedback connectivity can be analyzed by this method (i.e. $A \rightarrow B$ as well as $B \rightarrow A$ can be studied). We had planned to use the method of (Kaminski et al. 2001) for determining the significance of any potentially causal results, but the results were so clear (see *Results*) that a simple sign-test sufficed.

Finally, as an initial test regarding whether M1 spikes were driving mRN spikes, we performed cross-correlogram tests on all combination of M1 and mRN cells during the task period. We particularly looked for examples of statistically significant M1 peaks in firing (above background rate assuming a Poisson distribution, with a 99% confidence limit) just before an mRN cell fired (and vice versa, since the mRN projects back to M1 via the cerebellum and thalamus). The time window in which we searched for significant M1 spike timing was 25 ms before mRN firing up to 25 ms after mRN firing. To determine whether the total number of cases we found for each rat was above chance, we shuffled the bins of data across all possible combinations and counted the number of significant cross-correlogram peaks occurring randomly. We used a proportion test to determine whether the actual number significant cross-correlogram peaks was greater than chance, given the total number of shuffles for each rat.

3 Results

Histology indicated that the implants in all four rats were properly located, and produced a total of 132 M1 cells and 55 mRN cells in addition to usable LFPs from each area in each rat. Firing rates relative to task phases were first analyzed with perievent histograms. All task phases from the *Final Sniff* to the *Contact* phase had $\geq 10\%$ of cells in either or both areas modulated significantly (either upward or downward, at the 99% confidence level), and are shown in Fig. 4. Regarding our hypotheses about M1 driving mRN cells, the perievent histogram results led to two clear hypotheses to be addressed in the remaining analyses. First, many cells in the two recorded brain regions exhibited similar response profiles during the *Final Sniff* phase, and therefore, cells in one region (hypothetically, M1) may be either coherent with or driving cells in the other region. This hypothesis will be examined with directed coherence as well as with the spike cross-correlograms. Second, given the paucity of M1 cells firing above baseline during the *Slot-align* phase, the activity seen in the mRN during the *Slot-align* phase should *not* appear to be driven by M1 cells.

As our first test of the hypotheses stemming from the perievent histogram data, we performed directed coherence analyses on the M1 and mRN LFPs from 5 trials per rat, yielding a total of 20 analyzed trials. Figure 5(a) shows directional flow *from the M1 to the mRN*, whereas Fig. 5(b) shows directional flow *from the mRN to M1*. The bars at the far right of each graph relate the correlation coefficients to the color scheme of the graphs (with dark blue representing no directional flow and a correlation coefficient of 0). It is readily apparent that information flows from M1 to the mRN virtually continuously throughout the trial—even during the *Slot-align* phase, when a large percentage of mRN cells fired above baseline and M1 cells did not. Sixteen of the 20 analyzed trials showed two critical results (sign-test, $p < 0.02$): (1) Contrary to the perievent histogram-based hypothesis, the M1 appeared to be driving the mRN on the *Slot-align* phase, and (2) there was no evidence of the mRN driving M1 cells at any point in the trial. The extremely brief and weak appearance of the mRN driving M1 in Fig. 5(b) occurred approximately 100 ms before the *Slot-align* phase and lasted 10–20 ms at most. Furthermore, virtually none of the other trials displayed evidence of such an event at that time or indeed, at any other time. Finally, Fig. 5 displays one of the 4 outlier trials in which there was any sign, at any point in the trial, where mRN may have been driving M1 activity. Sixteen of the trials (80%) did not even show this meager level of ‘driving’ by the mRN. Thus, the throughout most or even all of the trial, including during the *Slot-align* behavioral event, directed coherence results indicated that information flowed from M1 to the mRN.

The above directed coherence results were based on analyses of local field potentials, which mainly reflect vector-summed dendritic activity but can also contain a significant spike component. Our next analysis more directly checked *whether M1 drove the spiking of mRN cells* through its own spike output, based on the assumption that spiking activity in the two areas would be correlated—with M1 spikes occurring slightly earlier than correlated mRN spiking—throughout the trial. For all recorded units within each animal, we checked for significant cross-correlogram peaks in all combinations of M1 and mRN units during the reaching trial periods only.

The results for one animal, with 29 M1 units and 23 mRN units, are presented in Fig. 6. The analysis revealed a small number of M1 cells firing above their background upon each mRN cell spike. Figure 6(a) shows three categories of putatively correlated firing. The first cross-correlogram (upper left corner) has a brief rise in firing of an M1 cell within 2 ms of the firing of the mRN cell, with a sparse background rate. The second graph in Fig. 6(a) (top middle) shows a similar temporal relationship between the M1 cell and the mRN cell, but with a much higher background firing rate. The third graph in Fig. 6(a) shows an oscillatory relationship, with the M1 cell’s firing rate rising and falling as mRN spikes occur, with a frequency of

roughly 20–30 Hz. Interestingly, very few of these oscillatory relationships were found within the reaching trial data, though they were noticeably more common when each animal's whole data file was examined (data not shown). Their oscillatory frequencies were far above chewing and grooming rates, and their significance is unclear.

But during the reaching trials, were there more pairs of M1 and mRN cells significantly cross-correlated than would be expected by chance? To compute this, for each pair of cells we shuffled bins of reach-period data in a fully combinatorial manner, and the number of significant cross-correlations occurring by chance was determined. Figure 6(b) shows the results for one animal. The number of significant cross-correlograms found in the animal's actual dataset was 10 out of 667 (29×23) pairs. When the data were shuffled across all possible shuffles for that rat, 18 such peaks occurred by chance. For this rat, therefore, nonsignificantly more cross-correlations were found in the randomized data than in the real, unshuffled data ($p > 0.15$). The results for the other three rats were similar, with chance-level proportions of possible M1-mRN cross-correlations (all p -values for these three rats > 0.35).

4 Discussion

In the experiments reported here, we used the known anatomy of the rat motor cortex and red nucleus to test our hypothesis that under the conventional interpretation of local field potential-based analyses of information flow, M1 would spuriously appear to drive (Sherman and Guillery 1998; Reichova and Sherman 2004) mRN neurons. M1 projections to the corticospinal tract send collaterals to mRN distal dendrites (Keifer and Lustig 2000), and therefore M1 dendritic activity or spikes would not be expected not to drive mRN spiking with large, long-distance traveling EPSPs. More likely, the M1 input would more subtly modulate the processing of driver inputs coming from the interpositus nucleus (Ralston 1994b; Tarnecki 2003; Tarnecki et al. 2001) or elsewhere. However, most scientists who use LFPs for information flow analyses conclude that if dendritic activity in area 1 (e.g., M1) precedes correlated dendritic activity in area 2 (e.g., the mRN), it can be concluded that area 1 drives the spiking of area 2 during that task.

The outcomes of our experiments and analyses demonstrate that such interpretations are easily misguided. First, we constructed perievent histograms centered around task phases known to be associated with significant modulation in one or both recorded regions (Hermer-Vazquez et al. 2004). These histograms allowed us to formulate two hypotheses—based on the conventional interpretation field potential-based information flow analyses—regarding spikes and drivers, on the one hand, and LFPs and modulators, on the other. Hypothesis 1 was that M1 and mRN spiking would be correlated, as revealed by cross-correlogram analysis, at least during the *Final Sniff* portion of the task. The perievent histograms revealed relatively large percentages of units in both areas with a similar response profile during that phase (a sharp firing rate decrease during a ~100 ms period centered on the *Final Sniff*). Although we analyzed each trial as a whole, some cross-correlation during this substantial portion of each trial (often > 100 ms out of a total trial length of ~350 ms) should have been evident if this hypothesis were true. Instead, *no* rat had any above-chance correlation between M1 and mRN spiking. Also, it would be expected that the Granger causality results might show M1 driving mRN cells during that phase, given the larger number of M1 cells displaying the characteristic *Final Sniff* response profile and the known anatomy of corticorubral networks. Yet despite this, the Granger causality analyses indicated that M1 was driving the mRN during this phase (and all the other phases). The most parsimonious interpretation of these data is that M1 projection cells modulated the distal dendrites throughout each trial, but did not drive mRN spiking.

Hypothesis 2 was that, for the only task phase in which many mRN neurons, but almost no M1 neurons, strongly fired above baseline—*Slot-align*, the Granger causality analyses should not

show the M1 ‘driving’ the mRN. Instead, what we found was that for 16 of 20 trials analyzed across 4 rats, the M1 appeared to ‘drive’ the mRN throughout the task, *even during this phase*. Furthermore, none of the 4 trials that showed any mRN driver activity whatsoever (as in Fig. 5) revealed any apparent mRN driving during the *Slot-align* time period. Thus, all 20 out of 20 analyzed trials showed that the M1 appeared to drive the spiking of mRN cells on *Slot-align*, even though virtually no M1 cells fired above baseline then while a large plurality of mRN cells did. Again, the most parsimonious interpretation of these results is that M1 projections modulated the excitability of mRN cells via their terminations on mRN distal dendrites, but did not actually drive the spike activity of mRN neurons.

These results are consistent with Sherman and colleagues’ distinction between driver and modulator activity (Sherman and Guillery 1998, 2002; Reichova and Sherman 2004), and they have strong implications for how analyses of information flow through brain networks are interpreted. M1’s terminations on mRN projection cell distal dendrites most likely leads to small EPSPs with a subtle influence on overall mRN activity. (It should be noted that in monkeys, some of the M1 projections also terminate on the perisomatic region of mRN inhibitory interneurons, with unknown functional consequences (Ralston 1994a).) Thus, LFP-based inferences about information flow may be best viewed within the modulator paradigm unless there is parallel spike evidence correlating area 1’s spikes with those in area 2. In contrast, research from other laboratories indicates that cerebellar interpositus projections onto mRN cells exert a much stronger influence on mRN spiking (Tarnecki 2003; Tarnecki et al. 2001; Thompson et al. 1997). Analyses of information flow based on spike data are currently more problematic because they require convolution of the discontinuous spike train data with a continuous kernel. Depending on one’s choice of kernel or the degree of noise in any of the analyzed time series, the preprocessing can introduce artifactual results (Albo et al. 2004; Ramsey and Lampart 1998; Sameshima and Baccala 1999; Baccala and Sameshima 2006). For example, spike data from different types of neuron (such as large-bodied pyramidal cells versus most interneurons) vary greatly in their signal-to-noise ratio, and cells (or other sources used in the analysis) with the greater signal can erroneously be found to be drivers (Albo et al. 2004).

However, the field of signal analysis is rich in new approaches for tracing information flow through complex neural networks (e.g. Bayesian network analysis; Pe’er 2005; Patel et al. 2006; Hartemink et al. 2001, 2002), although many of these techniques have not yet been adapted to the spectral domain. The current results demonstrate the dual needs for modelers to understand the neurobiological details of the networks they are studying, and the need for new techniques designed to handle different kinds of neural measures. For instance, to our knowledge there are currently no techniques that spectrally and temporally relate the spike output of a cortical projection layer (e.g. pyramidal cell output in cortical layer 5) to the dendritic processing at an efferent site. We are optimistic that such techniques will be developed, rigorously tested with neurally realistic surrogate data (e.g. times series that are nonstationary, non-Gaussian, highly transient fluctuations in simulated oscillatory LFP energy), and finally applied to real, N-by neural data streams.

Acknowledgements

The author thanks Mingzhou Ding, Luiz Baccala, Rodolfo Llinas, and Ray Hermer-Vazquez for useful discussions. This research was supported by a University of Florida Opportunity Seed Fund Award to L.H.-V.

References

- Albo Z, Di Prisco GV, Chen Y, et al. Is partial coherence a viable technique for identifying generators of neural oscillations? *Biol Cybern* 2004;90(5):318–326. [PubMed: 15221392]

- Baccala LA, Sameshima K. Partial directed coherence: a new concept in neural structure determination. *Biol Cybern* 2001;84(6):463–474. [PubMed: 11417058]
- Baccala LA, Sameshima K. Comments on ‘Is partial coherence a viable technique for identifying generators of neural oscillations?’: Why the term ‘Gersch causality’ is inappropriate: common neural structure inference pitfalls. *Biol Cybern* 2006;2:135–141. [PubMed: 16715246]
- Brovelli A, Ding M, Ledberg A, Chen Y, Nakamura R, Bressler SL. Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by Granger causality. *Proc Natl Acad Sci USA* 2004;101(26):9849–9854. [PubMed: 15210971]
- Bullock TH. Signals and signs in the nervous system: the dynamic anatomy of electrical activity is probably information-rich. *Proc Natl Acad Sci USA* 1997;94(1):1–6. [PubMed: 8990149]
- Chen Y, Bressler SL, Ding M. Frequency decomposition of conditional Granger causality and application to multivariate neural field potential data. *J Neurosci Methods* 2006;150(2):228–237. [PubMed: 16099512]
- Fanselow EE, Sameshima K, Baccala LA, Nicolelis MA. Thalamic bursting in rats during different awake behavioral states. *Proc Natl Acad Sci USA* 2001;98(26):15330–15335. [PubMed: 11752471]
- Geweke J. Measurement of linear dependence and feedback between multiple time series: rejoinder. *J Am Stat Assoc* 1982;77(378):323–324.
- Granger CWJ. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 1969;37(3):424–438.
- Gueguin M, Le Bouquin-Jeannes R, Faucon G, Chauvel P, Liegeois-Chauvel C. Evidence of functional connectivity between auditory cortical areas revealed by amplitude modulation sound processing. *Cereb Cortex* 2006;17(2):304–313. [PubMed: 16514106]
- Hartemink AJ, Gifford DK, Jaakkola TS, Young RA. Using graphical models and genomic expression data to statistically validate models of genetic regulatory networks. *Pac Symp Biocomput* 2001:422–433. [PubMed: 11262961]
- Hartemink AJ, Gifford DK, Jaakkola TS, Young RA. Combining location and expression data for principled discovery of genetic regulatory network models. *Pac Symp Biocomput* 2002:437–449. [PubMed: 11928497]
- Hermer-Vazquez L, Hermer-Vazquez R, Moxon KA, et al. Distinct temporal activity patterns in the rat M1 and red nucleus during skilled versus unskilled limb movement. *Behav Brain Res* 2004;150(1–2):93–107. [PubMed: 15033283]
- Kaminski M, Ding M, Truccolo WA, Bressler SL. Evaluating causal relations in neural systems: granger causality, directed transfer function and statistical assessment of significance. *Biol Cybern* 2001;85(2):145–157. [PubMed: 11508777]
- Keifer J, Lustig DG. Comparison of cortically and subcortically controlled motor systems, II: distribution of anterogradely labeled terminal boutons on intracellularly filled rubrospinal neurons in rat and turtle. *J Comp Neurol* 2000;416(1):101–111. [PubMed: 10578105]
- Kleim JA, Barbay S, Cooper NR, et al. Motor learning-dependent synaptogenesis is localized to functionally reorganized motor cortex. *Neurobiol Learn Mem* 2002;77(1):63–77. [PubMed: 11749086]
- Kus R, Kaminski M, Blinowska KJ. Determination of EEG activity propagation: pair-wise versus multichannel estimate. *IEEE Trans Biomed Eng* 2004;51(9):1501–1510. [PubMed: 15376498]
- Laughlin SB, Sejnowski TJ. Communication in neuronal networks. *Science* 2003;301(5641):1870–1874. [PubMed: 14512617]
- Mitzdorf U. Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol Rev* 1985;65(1):37–100. [PubMed: 3880898]
- Patel RS, Bowman FD, Rilling JK. A Bayesian approach to determining connectivity of the human brain. *Hum Brain Mapp* 2006;27(3):267–276. [PubMed: 16092131]
- Pe’er D. Bayesian network analysis of signaling networks: a primer. *Sci STKE* 2005;2005(281):pl4. [PubMed: 15855409]
- Ralston DD. Corticorubral synaptic organization in *Macaca fascicularis*: a study utilizing degeneration, anterograde transport of WGA-HRP, and combined immuno-GABA-gold technique and computer-assisted reconstruction. *J Comp Neurol* 1994a;350(4):657–673. [PubMed: 7534317]

- Ralston DD. Cerebellar terminations in the red nucleus of *Macaca fascicularis*: an electron-microscopic study utilizing the anterograde transport of WGA:HRP. *Somatosens Mot Res* 1994b;11(2):101–107. [PubMed: 7526581]
- Ramsey JB, Lampart C. Decomposition of economic relationships by timescale using wavelets: money and income. *Macroecon Dyn* 1998;2(1):49–71.
- Reichova I, Sherman SM. Somatosensory corticothalamic projections: distinguishing drivers from modulators. *J Neurophysiol* 2004;92(4):2185–2197. [PubMed: 15140908]
- Sameshima K, Baccala LA. Using partial directed coherence to describe neuronal ensemble interactions. *J Neurosci Methods* 1999;94(1):93–103. [PubMed: 10638818]
- Sharott A, Magill PJ, Bolam JP, Brown P. Directional analysis of coherent oscillatory field potentials in the cerebral cortex and basal ganglia of the rat. *J Physiol* 2005;562(Pt 3):951–963. [PubMed: 15550466]
- Sherman SM. Thalamic relays and cortical functioning. *Prog Brain Res* 2005;149:107–126. [PubMed: 16226580]
- Sherman SM, Guillery RW. On the actions that one nerve cell can have on another: distinguishing “drivers” from “modulators”. *Proc Natl Acad Sci USA* 1998;95(12):7121–7126. [PubMed: 9618549]
- Sherman SM, Guillery RW. The role of the thalamus in the flow of information to the cortex. *Philos Trans R Soc Lond B Biol Sci* 2002;357(1428):1695–1708. [PubMed: 12626004]
- Tarnecki R. Responses of the red nucleus neurons to limb stimulation after cerebellar lesions. *Cerebellum* 2003;2(2):96–100. [PubMed: 12880176]
- Tarnecki R, Lupa K, Niechaj A. Responses of the red nucleus neurons to stimulation of the paw pads of forelimbs before and after cerebellar lesions. *J Physiol Pharmacol* 2001;52(3):423–436. [PubMed: 11596860]
- Thompson RF, Bao S, Chen L, et al. Associative learning. *Int Rev Neurobiol* 1997;41:151–189. [PubMed: 9378587]
- Whishaw IQ, Pellis SM. The structure of skilled forelimb reaching in the rat: a proximally driven movement with a single distal rotatory component. *Behav Brain Res* 1990;41(1):49–59. [PubMed: 2073355]

Fig. 1.

Possible misleading interpretations of Granger causality analyses, particularly when local field potentials or EEG signals are used. **a** Granger causality analysis where area 1 appears to drive area 2. However, if only LFPs are used for the analysis, it is not safe to assume that dendritic activity in area 1 is 'driving' any form of activity in area 2. **b** An often-cited interpretive problem with Granger causality analyses, in which information spuriously appears to flow from area 1 to area 2. However, when conditional Granger analysis is applied to a third site that receives inputs from area 1 and projects to area 2, analysis of the three time series reveals that information flows from site 1 to site 3 and finally to site 2, a different causal chain. **c** The interpretative problem dealt with in this paper: If only dendritic or modulatory activity is

analyzed for two or more time series, area 1 might appear to send information to area 2 when in fact areas 2 and 3 are simultaneously active and only area 3 is sending information to area 2

Fig. 2.

Recording sites and their neuroanatomy in the rat. We recorded from the caudal forelimb motor cortex (M1) and the magnocellular red nucleus. M1 projection cells form part of the corticospinal tract, and send collaterals to the distal dendrites of the red nucleus. In contrast, output cells in the interpositus nucleus of the cerebellum (ipCB) project onto or near the soma of red nucleus cells (Keifer and Lustig 2000)

Fig. 3.

a Depiction of a trained rat reaching through a narrow plastic slot to obtain a food pellet in one of two pellet wells on a shelf in front of him. Each of the two pellet wells is 13 mm from the edge of the slot. **b** Time course of the average single reach in this study. On a typical trial, the rat places its nose on the shelf as close to the pellet as possible, sniffs for 1–3 respiratory cycles, lifts its paw and grasps the pellet. Rats in this study were trained to their own asymptotic levels, with a final group average performance for grasping and eating the pellet of ~70%

Skilled reaching

TOP: M1 unit subpopulations

Fig. 4.

Perievent histograms of cell firing in each recorded area centered around the major behavioral events of each trial. Percentages of cells active above or below baseline (upper and lower 99th percentiles provided by the dotted lines) are depicted when $\geq 10\%$ of units in each recorded area fell into a given category (such as above-baseline activity as the reaching paw was lifted). The binwidth was 25 ms. Both M1 and the mRN show relatively large proportions of cells with an inhibited firing rate during the *Final Sniff* before lifting the paw. However, only the M1 recordings yielded many cells active above baseline on the *Lift* of the paw, the moment when the paw first touches the shelf (*Onshelf*), and the time of the paw contacting the pellet (*Contact*). In contrast, many mRN cells, and virtually no M1 cells, were active on the moment at which the upper trunk, elbow and wrist were simultaneously moved to begin projecting the paw through the slot (*Slot-align*). From these data, we hypothesized that during the *Slot-align* phase, M1 spikes would not be found to “drive” or even correlate with mRN spikes

Directed coherence between M1 and the mRN on one representative reaching trial

Fig. 5.

Directed coherence between M1 and the mRN on one of the 4 trials that displayed *any* mRN 'driving;' 16 out of the 20 trials displayed no apparent driving whatsoever. The analysis was performed on local field potentials recorded from pairs of M1 and mRN wires on a trial-by-trial basis, with each rat providing 5 trials. The trials were randomly selected except that trials with railing LFPs were omitted and randomly replaced. The top rectangular graph (A) depicts flow *from M1* to the mRN, and the bottom graph (B) depicts flow *from the mRN* to M1. For each graph, trial time is in ms on the x-axis, frequency from 0–100 Hz is on the y-axis, and energy intensity is color coded, with bright red for highest energy and dark blue for lowest energy. The fact that each pair of analyzed signals were normalized to each other by making energy in the theta band 0 dB and that there was substantial classical coherence between the two signals (data not shown; Hermer-Vazquez et al., in prep.) suggests that a lack of energy in the mRN did not account for the nearly unidirectional flow throughout the trial. Nearly all analyzed trials (16/20) showed this pattern of results

Fig. 6.

Cross-correlations between M1 and mRN unit activity during reaching trials. This analysis examined whether an above-chance number of M1 unit spikes preceded an mRN spike by 0–10 ms, which would suggest a “driving” role for at least some M1 units. **a** The x-axes show time relative to each mRN cell spike, from 100 ms before that event until 100 ms afterward. The y-axes shows firing intensity, with dotted lines marking the upper 99% confidence limit. The first two graphs show examples of M1-mRN unit pairs that were significantly cross-correlated during one rat’s reaches, with one M1 cell spiking slightly before the mRN cell in the first graph (*upper left*), and another M1 cell firing nearly coincidentally with an mRN cell in the upper middle graph. The upper right graph shows an oscillatory relationship between M1 and mRN firing during the reaching trials, with a cycle of ~25 Hz. Very few (<10) examples of this type were found across the trials from the four rats. The oscillatory cross-correlograms were smoothed over 3 bins to make the oscillatory dynamics clearer. **b** Test of the hypothesis that an above-chance number of M1 spikes would precede or occur nearly coincidentally with mRN spikes during reaching for each rat. On a unit pair by unit pair (across the M1 and mRN) basis, bins of data from all reaching trials for each animal were shuffled in a fully combinatorial manner, and the number of significant cross-correlations occurring by chance was determined. Results for one animal are shown here. The number of significant cross-correlograms found in the animal’s actual dataset was 10 out of 667 pairs. When the data were shuffled across all possible shuffles for that rat, 18 such peaks occurred—by chance. Thus, a *nonsignificantly greater* number of cross-correlations was found in the randomized data than in the real, unshuffled data