Published in final edited form as:

J Neurosci Methods. 2014 January 30; 222: 24–33. doi:10.1016/j.jneumeth.2013.10.005.

Short-time windowed covariance: A metric for identifying non-stationary, event-related covariant cortical sites

Timothy Blakely^{a,*}, Jeffrey G. Ojemann^b, and Rajesh P.N. Rao^c

^a Department of Bioengineering, University of Washington, United States

^b Department of Neurosurgery, University of Washington, United States

^c Department of Computer Science, University of Washington, United States

Abstract

Background—Electrocorticography (ECoG) signals can provide high spatio-temporal resolution and high signal to noise ratio recordings of local neural activity from the surface of the brain. Previous studies have shown that broad-band, spatially focal, high-frequency increases in ECoG signals are highly correlated with movement and other cognitive tasks and can be volitionally modulated. However, significant additional information may be present in inter-electrode interactions, but adding additional higher order inter-electrode interactions can be impractical from a computational aspect, if not impossible.

New method—In this paper we present a new method of calculating high frequency interactions between electrodes called Short-Time Windowed Covariance (STWC) that builds on mathematical techniques currently used in neural signal analysis, along with an implementation that accelerates the algorithm by orders of magnitude by leveraging commodity, off-the-shelf graphics processing unit (GPU) hardware.

Results—Using the hardware-accelerated implementation of STWC, we identify many types of event-related inter-electrode interactions from human ECoG recordings on global and local scales that have not been identified by previous methods. Unique temporal patterns are observed for digit flexion in both low- (10 mm spacing) and high-resolution (3 mm spacing) electrode arrays.

Comparison with existing methods—Covariance is a commonly used metric for identifying correlated signals, but the standard covariance calculations do not allow for temporally varying covariance. In contrast STWC allows and identifies event-driven changes in covariance without identifying spurious noise correlations. Conclusions: STWC can be used to identify event-related neural interactions whose high computational load is well suited to GPU capabilities.

Keywords

Electrocorticography; Covariance; GPU; Acceleration

1. Introduction

Electrocorticographical (ECoG) recordings have provided the computational neuroscience community with unprecedented spatial and temporal resolution of brain activity. Because of

the high signal-to-noise ratio and single-trial fidelity, ECoG is increasingly explored in a variety of tasks, including motor function. ECoG has been shown to be a viable control signal for use in brain-machine interfaces (Pfurtscheller et al., 2000; Wolpaw et al., 2003; Miller et al., 2008; Blakely et al., 2009; Brunner et al., 2011; Leuthardt et al., 2011) and is also used during cortical mapping (Crone et al., 1998, 2006; Sinai et al., 2005; Hermes et al., 2012; Wray et al., 2012). Though the neuroanatomy underlying the signals recorded by ECoG are up for discussion (see Crone et al., 1998; Sarnthein et al., 2003; Miller et al., 2007), there are measurable characteristic changes that are highly correlated with local neural activity nonetheless.

Characteristic spectral changes have been identified that are robust across patients, experimental methods, and brain areas. Decreases in low-frequency (“beta”, “mu” and “alpha” ranges) are seen over large areas of cortex during motor movement, while broad-band high frequency increases occur in spatially focal areas that have been previously associated with their respective movement (Sinai et al., 2005; Crone et al., 2006; Miller et al., 2007, 2009, 2010; Blakely et al., 2008).

Because the brain is highly connected and these broad-band increases are not limited to one area of the brain (Darvas et al., 2004; Anderson et al., 2013; Wander et al., 2013), it is only natural to assume that areas of the brain interact and communicate information.

One of the most prominent issues with exploring multi-electrode interaction is the issue of dimensionality explosion. Calculating a metric for any permutation of two channels is a large space to explore, and adding additional parameters to explore can increase the number of calculations by an order of magnitude or more. For example, checking 5 frequency bands across pairs in a total of 64 electrodes requires 20,480 separate calculations. If the time required for each calculation is even as short as a single second, performing such a calculation would take upwards of 6 h in total.

This effect is compounded when looking for non-stationary, event-related correlations. The human brain operates in a highly redundant, parallel fashion; action potentials and synaptic neurotransmitter release occurs in nearly all regions of the brain. That does not exclude forms of cortical synchronization such as thalamocortical rhythms, commonly implicated in motor control (Sinai et al., 2005; Miller et al., 2009, 2010), but it is entirely possible that areas of the brain can initiate transfer of information from one area to a spatially remote area via a transiently-activated network, only to cease communication over this pathway once the information transfer is complete. To identify these event-related interactions, it is necessary to do finer-scale calculations by dividing the data up into smaller sections. Many algorithms make the assumption that brain areas are continuously correlated due to computational constraints, which would likely not illuminate short timescale event-related correlations.

We present a new method of exploring high dimensional cortical data under the assumption that event-related high frequency increases in one area of cortex are closely correlated with a remote area for only short periods of time. Short-time windowed covariance (STWC) is a measurement of lagged covariance of two signals over a short time period. Calculating the covariance of two recorded signals implies that the two recorded signals are continuously correlated during that time period. STWC, as its name implies, calculates the covariance over a short window lagged by $\pm \delta$ samples. This window is recalculated along every time point recorded for every pair of channels, across a number of different lags (δ). Exploring such a high dimensional space has a high computational load associated with it; to solve that problem, we provide an implementation of the algorithm on a consumer-grade graphics processing unit (GPU) commonly found in desktop computers today. Repurposing GPU hardware (originally designed to do streaming manipulation of geometric primitives) allows

us to take advantage of the massively parallel architecture available within the chipset, greatly speeding up the algorithm. By leveraging hardware acceleration, exploring high dimensional spaces with STWC is no longer impractical nor impossible.

2. Method

2.1. Short-time windowed covariance (STWC)

STWC calculates covariance measures (with the option to normalize to correlation) over two signals at every time point. A covariance calculation is performed with a small sample window of the source channel centered on sample t with multiple windows of the remote channel:

$$Q(x, y, t, \tau, \delta) = \sum_{i=t-\frac{\tau}{2}}^{t+\frac{\tau}{2}} \frac{(x_i - \bar{x})(y_{\delta+i} - \bar{y}_{\delta})}{\tau+1} \quad (1)$$

where x and y are the source and destination channels, τ is the window size, and δ is the number of samples the window by which y is shifted. This calculation is performed over the permutation of possible channels x and y , with each pair iterating through all possible samples of t , for each lagged value δ from the desired lag window of $[-\Delta, \Delta]$ (see Fig. 1). It is easy to see that while the complexity of an individual calculation is low, the number of possible combinations of parameters can be extraordinarily large.

The selected windows in the second channel are the same length of the source channel's window but shifted from $-\delta$ to δ , giving $2 \times \delta + 1$ covariance measures for each sample point t .

While at first glance it would appear that cross-covariance and STWC would give similar answers given that the method deals with lagged windows of data, they answer fundamentally different questions. Cross-covariance takes time-locked signals at the same time point and shifts them forward and backward in time, asking the question “what is the covariance between signals x and y , taken from the same sample, if we shift them in time?”. In contrast, STWC provides a measurement of how covariant two signals are over time when one is lagged by δ samples, asking “what is the covariance between two windows from x and y when a window of data is taken at time point t and one is taken at sample $t + \delta$, over the range $\delta = [-\Delta, \Delta]$?” By keeping the two signals being compared at a fixed lag and computing over each sample, STWC computes correlated activity in one channel relative to another on a given timescale and allows the strength of this covariance to change over time.

One of the largest advantages to this measurement is that covariance is a well-understood and widely known mathematical technique that has been widely used in previous neuroscience studies (Crone et al., 1998; Mechelli et al., 2005; Wang et al., 2010). STWC builds on covariance by calculating values on short segments of data allowing the covariance to change over time. This is the key factor when looking for remote interactions of correlated neural activity: brain-regions that are not normally covariant but become correlated only around an event that requires coordination between two or more areas of cortex.

2.2. CUDA

STWC is calculated in such a way that each calculation of parameters is independent for all values of t , τ and δ . Therefore each covariance calculation can be computed simultaneously and in parallel, a form of parallelism known as *embarrassingly parallel*. In contrast to

modern computer processors that have two or four cores – units capable of simultaneous computation – modern graphics processor units (GPUs) have hundreds of cores, giving them the capability to perform embarrassingly parallel tasks orders of magnitude faster. Though their core count often comes with tradeoffs like limited memory and the inability to change programs mid-execution, they are particularly well suited to perform vector calculations on large amounts of input data; exactly the type of computation required by STWC.

We have implemented STWC using the CUDA mathematical library by nVidia, a functional library that allows the program to execute on a computer's graphics processing unit (GPU) instead of the core CPU allowing for massively parallel computation to occur. The STWC library takes in two vectors of single-precision floating point values and a desired lag-window size and divides the computation up into individual calculations of Eq. (1), optimized around the available GPU hardware. Depending on the number of cores available and maximum size of the computation, it allocates blocks of computations known as *threads* in order to maximize the usage of available processors. Current GPU standards are capable of executing 32 threads on a single arithmetic core, known as a *warp*. Maximizing the number of threads filled in each warp ensures that the algorithm is running at maximum performance. Warp occupancy and algorithm decomposition is illustrated in Fig. 3.

In order to show the performance of our custom GPU-accelerated algorithm, we tested it against an implementation of the algorithm written in C++ and running on a contemporary CPU (see Hardware, below). For uniformity, both the CPU and GPU versions of our custom STWC algorithm were compiled and wrapped as a Matlab plugin that was callable from the Matlab environment.

2.3. Parallel implementation

To compute STWC on a pair of channels requires calculating covariance measures at N different lags for every time point T . These dimensions result in a total of $N \times T$ covariance calculations, which can be viewed as a matrix of $N \times T$ covariance measures. As each of the individual covariance calculations are inherently independent, we can process each covariance measure in parallel. Fig. 3 shows how the algorithm maps each covariance computation to an individual CUDA thread. The computation is initially divided along the lag dimension: if the number of lags is less than the GPU's upper limit on block size (512 threads per block for test hardware) a single kernel execution is performed; otherwise the data spread among the minimum number of kernel executions required ($\text{floor}(N/512) + 1$). Subsequent division of the data into thread blocks was designed to maximize warp efficiency and minimize idle threads.

Memory access was optimized to leverage block-level shared memory. At any given time point t , the full window of the stationary channel and all possible values of the lagged channel were read into block-level shared memory, with CUDA threads synchronizing after each step. Because this memory read was pulled into block-shared memory, each block only required this read to be performed exactly once for every time point, greatly lowering memory access times compared with having each CUDA thread reading from global memory independently.

2.4. Artificial data

While STWC is developed to identify covariant changes in high frequency broadband signals, it must not identify spurious covariance measures when applied to cortical ECoG recordings. Two simulated cortical traces were generated (called *source* and *target*) much like those commonly recorded from human ECoG studies where the power on each frequency falls off at approximately $1/f^\alpha$.

Each signal was the summation of random sinusoids with frequencies between 0.1 Hz and 200 Hz and a total length of two minutes, assuming a sampling rate of 1200 Hz, according to:

$$u(t) = \sum_{f=1}^{200} \frac{\sin(2\pi ft + \psi)}{f^x} \quad (2)$$

where ψ is a randomly generated phase lag based on sample t . The sinusoids were varied in the phase domain and were assumed to fall off at a rate of $1/f^x$ (Miller et al., 2009). To capture the high frequency, each signal was filtered with a 4th order butterworth filter from 75 to 200 Hz. A PSD was calculated on each signal to ensure that their power spectra were comparable to those seen in humans (see Fig. 2).

Two simulations were performed. In the first dataset, random uncorrelated increases in log power in the broadband high frequency range were added to both the source and target channels at random intervals, simulating two channels that show relative increases in high gamma, yet not correlated in time. Each increase lasted 100 ms and was modeled by multiplying the log power of each channel (obtained by applying a Hilbert transform after the butterworth) by a gaussian scaling factor of 1.1. Next, a second dataset was generated by taking one of the original simulated channels and applying 30–50 ms-long gaussian increases in log high frequency power, akin to the first dataset. At each increase in the source channel, a corresponding increase was added to the destination channel exactly 50 ms later.

STWC was run on both datasets, with a lag window of size $\delta = 200$ (165 ms at a sampling rate of 1.2 kHz). At each point t where an increase in high gamma was introduced into the source channel, a window of ± 500 ms of data was taken and averaged over the 30 instances (Fig. 3).

2.5. Experimental design and validation

To assess the algorithms' performance on real-world cortical data, we designed an experiment to capture stereotyped, coordinated hand motor movements and time-lock them to cortical signals recorded via electrocorticography. Two neural recordings were performed in two separate subjects implanted with subdural electrocorticographic disk electrodes (see Fig. 4). Subjects were presented with a computer screen situated comfortably in visual range and were asked to don a 22 degree-of-freedom dataglove (CyberGlove Systems, Inc.) that records joint angles in nearly all the joints in the hand. A blank silhouette of a hand is shown to the subject during periods of rest. At the start of a single trial, one of the five fingers would be highlighted in a unique color, or the thumb and index finger would be highlighted the same color to signify a "pinch". Subjects were instructed to flex the same finger that was highlighted on the screen three times (or make a pinching motion for pinch), then relax (see Fig. 5). After 2000 ms, the highlight is removed and the on-screen hand silhouette becomes blank again. This was repeated 10 times for each finger and an additional 10 times for pinch, with random presentation order.

Subject A underwent a standard long-term implantation of an ECoG array for the surgical treatment of intractable epilepsy. The electrode array used in subject A was an 8×8 subdural platinum electrode array (Ad-Tech, Racine, WI) covering primary motor cortex. Electrodes on this grid were 3 mm in diameter with 2.3 mm exposed surface diameter and spaced at 10 mm center-to-center, covering a total of. The experiment was carried out 4 days after implantation at the patient's bedside by splicing the clinical system's recording leads into a g.USBamp biosignal amplifier (G.Tec, Austria) and sampling at 1.2 kHz. Neural

signals and dataglove samples were collected and time-locked by the BCI2000 recording suite (Schalk et al., 2004).

Subject B was recorded using a higher resolution ECoG array; a grid of 8×8 platinum electrodes 1.5 mm in diameter (1.1 mm exposed) and spaced at 3 mm center-to-center. This grid was placed over primary hand motor cortex during an awake craniotomy and the experiment subsequently performed in the operating room. Aside from the grid size, the experimental protocol was identical to that of subject A.

Both subjects gave informed written consent in accordance with the University of Washington IRB and HIPAA protocols.

2.6. Hardware

CPU benchmarks were performed on an Intel Core i7-2720QM CPU (4-cores) operating at 2.20 GHz. C++ code was compiled using Visual Studio 2010 using optimized compiler flags under release mode. CUDA code was run on an nVidia GeForce GTX 580M graphics card, compiled by the nVidia NVCC compiler using CUDA runtime version 4.0. Benchmarks were performed using 1, 2, and 4 threads on the CPU, and CUDA compute capability 2.1 for GPU. Note that the GPU's arrangement of cores was 8 streaming multiprocessors, each containing 48 CUDA cores. In both cases the operating system was Windows 7.

3. Results

3.1. Generated data

When STWC is applied to the uncorrelated simulated ECoG data, no increases or decreases in covariance are observed (Fig. 2c). This is consistent with the expected result of two random signals having low covariance measurements. When the increases in high gamma were introduced (Fig. 2d), the resulting STWC output showed a high increase in covariance in the second channel that was consistently delayed by 50 ms. Periods before and after the increase showed no changes in covariance. These results demonstrate STWC does not identify spurious covariance, is capable of identifying specific delays in covariance, and can identify event-related increases in covariance.

3.2. Event-related temporal information

Applying STWC to the human ECoG data recorded in subject B shows clear differences in temporal covariance. By averaging over the peak of index finger flexion, clear increases in covariance occur in different pairs of electrodes, and in some cases even pairs of electrodes sharing the same source electrode (Fig. 6). Coincident increases in covariance occur during motor planning and initiation (Fig. 6d) in one pair of electrodes, while leading and lagged increases in covariance occur prior and during motor execution (Fig. 6c, e) at -53 ms and 98 ms, respectively. In addition, areas outside of index finger flexion show no increases or decreases in covariance, with any increases being temporally brief and not continuous pre, during and post motor execution. This demonstrates that the increases in covariance illustrated by STWC are nonstationary event-related increases.

3.3. Event-specific interactions

While it is useful to know that the increases in covariance are event-related, the patterns observed are unique across events as well. In the data recorded from Subject A, unique patterns of activity can be seen across flexions of the five different fingers (Fig. 7). Increases in covariance occur in pairs of electrodes predicted by anatomical coverage, including motor-sensory and sensory-premotor. However, areas that are not expected to interact significantly during motor execution (i.e. superior temporal and parietal) show no change in

covariance. It is important to note that the inter-electrode spacing is larger for Subject A than for Subject B, implying that STWC is viable for finding covariance inter-electrode increases in high gamma across a range of electrode sizes and spacing.

3.4. Global event-specific patterns

Global covariance patterns across all electrodes on the grid are unique across all event types. At a macro scale, Subject A showed areas of high covariance in sensorimotor areas, with additional increases seen in premotor and frontal areas (Fig. 8). At a more local level, patterns in covariance changes can be seen in Subject B's microgrid (Fig. 9). A number of electrodes show very high digit specificity for increased covariance, each being spatially distributed across the coverage of the grid.

3.5. Benchmarks

Benchmarking was performed by performing STWC on 100,000 samples of random single-precision floating-point data at 1200 Hz sampling rate, with window size of 300 samples and pre-/post-lag of 300 samples. CPU benchmarks were performed with 1, 2, 3 and 4 threads (maximum number of cores available). Table 1 shows the results of the benchmark, with the GPU showing a 30.9× speedup compared to the single-threaded algorithm and an improvement of 31.52× relative to the 4-thread version. It should be noted that even though additional worker threads are used, the relative speedup is not linear. Moving from one to two threads results in a speedup of 59%, the gains diminish as more threads are added (74% and 79%, respectively). This suggests that the bottleneck for the CPU is not raw computational power but available memory bandwidth.

During CUDA STWC computation, GPU warp occupancy was 82.7%, compared with the theoretical limit of 83.3%, according to the benchmarking toolset Parallel Nsight (nVidia Corporation), indicating that the thread allocation algorithm is operating at near optimal limits.

4. Discussion

A common problem when working with any type of biological data is the inherently low signal-to-noise ratio. One of the most common ways to combat this is to oversample the data, be it multiple experimental runs, higher sampling frequency in hardware or other means of generating additional data points. However, this additional data presents a challenge during analysis: a significant increase in post-experiment computational load. The additional data points could illuminate a new finding, but due to the computational time cost most such analysis is merely alluded to in discussion sections. Combine the additional complexity with the combinatorial problem of searching a high-dimensional space and one can easily see how computationally constrained biological data analysis truly is. Other studies have begun to leverage commodity hardware to accelerate neuroscientific analysis, including feature-extraction during brain-computer interfaces (Wilson and Williams, 2009) and cortical surface reconstruction from MRI images (Fischl, 2012).

In developing STWC, we were looking for a simple measure that could tease out event related covariance changes in high gamma—a measure commonly associated with local increases in noise-like neural activity. To do so required taking covariance measurements at as many time points as possible across multiple combinations of channels. We quickly realized that conventional computational methods and tools would not suffice and instead chose to implement it on a GPU that provides massive parallelism and computational horsepower in a consumer-grade device.

STWC itself is an easy measurement to understand from a mathematical perspective, producing a measure of covariance that can change over time. Initial presentation of the results can identify areas that interact in a transient way, but further extensions of the data may shine light on the temporal characteristic of inter-area information transfer around the cortex. By providing continuous temporal hints, the output of STWC could be used by machine learning and regression algorithms like Gaussian processes (Wang et al., 2010) or Granger causality (Winterhalder et al., 2005; Korzeniewska et al., 2008; Wilke et al., 2009) to concretely identify the timing and generalized information flow around the cortex. The scope of such studies is outside the scope of this paper, but we are currently investigating STWC in multiple experimental contexts with exciting results.

We hope this algorithm and its findings can provide some insight into the utility of consumer-grade GPU hardware for solving the computational bottlenecks facing the neuroscience community. To that end, the framework we have developed is not specific to STWC. We hope that by releasing it as an open-source library that other computationally constrained algorithms such as Granger causality and biphasic-coherence (Darvas et al., 2009) can benefit from GPU-grade computational horsepower.

Acknowledgments

Research reported in this publication was supported by the National Institute of Neurological Diseases of the National Institute of Health under award number R01NS065186, NSF award no. 0930908 and EEC-1028725, and the Keck Foundation. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

- Anderson N, Blakely T, Brunner P, Krusienski D, Moran D, Leuthardt E. High-frequency spectral changes in dorsolateral prefrontal cortex for potential neuroprosthetics. IEEE-EMBC. 2013
- Blakely, T.; Miller, KJ.; Rao, RPN.; Holmes, MD.; Ojemann, JG. Localization and classification of phonemes using high spatial resolution electrocorticography (ECoG) grids.. Conference proceedings: Annual International Conference of the IEEE Engineering in Medicine and Biology Society IEEE Engineering in Medicine and Biology Society Conference; 2008; p. 4964-7. Available at: http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4650328
- Blakely T, Miller KJ, Zanos SP, Rao RPN, Ojemann JG. Robust, long-term control of an electrocorticographic brain-computer interface with fixed parameters. Neurosurgical Focus. 2009; 27:E13. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/19569888>. [PubMed: 19569888]
- Brunner P, Ritaccio AL, Emrich JF, Bischof H, Schalk G. Rapid Communication with a “P300” Matrix Speller Using Electrocorticographic Signals (ECoG). Frontiers in Neuroscience. 2011; 5:5. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3037528&tool=pmcentrez&rendertype=abstract>. [PubMed: 21369351]
- Crone NE, Miglioretti DL, Gordon B, Lesser RP. Functional mapping of human sensor-motor cortex with electrocorticographic spectral analysis. II. Event-related synchronization in the gamma band. Brain. 1998; 121:2301–15. [PubMed: 9874481]
- Crone NE, Sinai A, Korzeniewska A. High-frequency gamma oscillations and human brain mapping with electrocorticography. Progress in Brain Research. 2006; 159:275. [PubMed: 17071238]
- Darvas F, Miller KJ, Rao RPN, Ojemann JG. Nonlinear phase-phase cross-frequency coupling mediates communication between distant sites in human neocortex. Journal of Neuroscience. 2009; 29:426–35. [PubMed: 19144842]
- Darvas F, Pantazis D, Kucukaltun-Yildirim E, Leahy RM. Mapping human brain function with MEG and EEG: methods and validation. Neuroimage. 2004; 23:S289. [PubMed: 15501098]
- Fischl B. FreeSurfer. Neuroimage. 2012
- Hermes D, Miller KJ, Vansteensel MJ, Aarnoutse EJ, Leijten FSS, Ramsey NF. Neurophysiologic correlates of fMRI in human motor cortex. Human Brain Mapping. 2012; 33:1689–99. [PubMed: 21692146]

- Korzeniewska A, Crainiceanu CM, Ku R, Franaszczuk PJ, Crone NE. Dynamics of event-related causality in brain electrical activity. *Human Brain Mapping*. 2008; 29:1170–92. [PubMed: 17712784]
- Leuthardt EC, Gaona C, Sharma M, Szrama N, Roland J, Freudenberg Z, et al. Using the electrocorticographic speech network to control a brain–computer interface in humans. *Journal of Neural Engineering*. 2011; 8:036004. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/21471638>. [PubMed: 21471638]
- Mechelli A, Friston KJ, Frackowiak RS, Price CJ. Structural covariance in the human cortex. *Journal of Neuroscience*. 2005; 25:8303–10. [PubMed: 16148238]
- Miller, KJ.; Blakely, T.; Schalk, G.; Den Nijs, M.; Rao, RPN.; Ojemann, JG. Three cases of feature correlation in an electrocorticographic BCI. Conference proceedings: Annual International Conference of the IEEE Engineering in Medicine and Biology Society IEEE Engineering in Medicine and Biology Society Conference; 2008; p. 5318-21. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/19163918>
- Miller KJ, Hermes D, Honey CJ, Sharma M, Rao RPN, Den Nijs M, et al. Dynamic modulation of local population activity by rhythm phase in human occipital cortex during a visual search task. *Frontiers in Human Neuroscience*. 2010; 4:197. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2990655&tool=pmcentrez&rendertype=abstract>. [PubMed: 21119778]
- Miller KJ, Schalk GS, Rao RPN, Leuthardt EC, Moran DW, Dennijs M, et al. Cortical spectral changes during actual and imagined motor movement, and the augmentation of spectral change with feedback. *Society for Neuroscience Abstracts Online*. 2007
- Miller KJ, Sorensen LB, Ojemann JG, Den Nijs M. Power-law scaling in the brain surface electric potential. *PLoS Computational Biology*. 2009; 5:e1000609. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2787015&tool=pmcentrez&rendertype=abstract>. [PubMed: 20019800]
- Pfurtscheller G, Neuper C, Guger C, Harkam W, Ramoser H, Schlogl A, et al. Current trends in Graz brain–computer interface (BCI) research. *IEEE Transactions on Rehabilitation Engineering*. 2000; 8:216–9. [PubMed: 10896192]
- Sarnthein J, Morel A, Von Stein A, Jeanmonod D. Thalamic theta field potentials and EEG: high thalamocortical coherence in patients with neurogenic pain, epilepsy and movement disorders. *Thalamus & Related Systems*. 2003; 2:231–8.
- Schalk G, McFarland DJ, Hinterberger T, Birbaumer N, Wolpaw JR. BCI2000: a general-purpose brain–computer interface (BCI) system. *IEEE Transactions on Bio-medical Engineering*. 2004; 51:1034–43. Available at: http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=1300799. [PubMed: 15188875]
- Sinai A, Bowers CW, Crainiceanu CM, Boatman D, Gordon B, Lesser RP, et al. Electrocorticographic high gamma activity versus electrical cortical stimulation mapping of naming. *Brain*. 2005; 128:1556–70. [PubMed: 15817517]
- Wander, Jeremiah D, Blakely T, Miller KJ, Weaver K, Johnson L, Olson J, et al. Distributed cortical adaptation during learning of a brain–computer interface task. *Proceedings of the National Academy of Sciences*. 2013
- Wang, Z.; Ji, Q.; Miller, KJ.; Schalk, G. Decoding finger flexion from electro-corticographic signals using a sparse Gaussian process.. 2010 20th International conference on pattern recognition; 2010; p. 3756-9. Available at: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5597572>
- Wilke C, Drongelen W, van Kohnen M, He B. Identification of epileptogenic foci from causal analysis of ECoG interictal spike activity. *Clinical Neurophysiology*. 2009; 120:1449–56. [PubMed: 19616474]
- Wilson JA, Williams JC. Massively parallel signal processing using the graphics processing unit for real-time brain–computer interface feature extraction. *Frontiers in Neuroengineering*. 2009; 2. [PubMed: 19277218]
- Winterhalder M, Schelter B, Hesse W, Schwab K, Leistriz L, Klan D, et al. Comparison of linear signal processing techniques to infer directed interactions in multivariate neural systems. *Signal Processing*. 2005; 85:2137–60.

- Wolpaw JR, McFarland DJ, Vaughan TM, Schalk G. The Wadsworth Center brain–computer interface (BCI) research and development program. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2003; 11:204–7. Available at: <http://ieeexplore.ieee.org/xpl/freeabs.all.jsp?arnumber=1214722>. [PubMed: 12899275]
- Wray CD, Blakely TM, Poliachik SL, Poliakov A, McDaniel SS, Novotny EJ, et al. Multimodality localization of the sensorimotor cortex in pediatric patients undergoing epilepsy surgery. *Journal of Neurosurgery: Pediatrics*. 2012; 10:1–6. [PubMed: 22681317]

HIGHLIGHTS

- Short-time windowed covariance (STWC): a metric quantifying temporal relationships of cortical activity.
- STWC shows event-related increases in cortical signal covariance.
- Applied STWC to human electrocorticographic (ECoG) signals recorded from the cortical surface.
- Demonstrated temporal directionality in high frequency ECoG during finger flexion.
- A software implementation of STWC that uses consumer graphics processing units (GPUs).

Fig. 1.

Overview of short-time windowed covariance. (A) Two of the recorded ECoG electrodes are chosen and band-passed for high frequency power, a measure indicative of local cortical activity. (B) At every time point, a window of data is chosen from the source and destination channels (blue and red trace resp.), with the destination window being the size of the source window plus $2 \times \delta$, where δ is the maximum time lag to compare. (C) For each lag $[-\delta, +\delta]$, the covariance is calculated between the source channel and the destination channel at $t \pm \delta$. (D) This gives the short-term covariance at time t if you lag the destination channel by δ . (E) This calculation is performed at every time point in the signal, resulting in a continuous time covariance that can illuminate nonstationary covariance in high frequency signals that could be associated with events like finger flexion (note: simulated data in E). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Fig. 2.

(A) Random signals were generated that contained the same basic traits as human ECoG recordings – namely a falloff in power of $1/f^x$. (B) Power spectra for the two signals band-passed from 75 to 200 Hz. (C) Uncorrelated increases in high gamma (as illustrated by the example traces in blue and red below) do not show any specific event-related covariance when centered on increases in the source channel. (D) However, increases in covariance can clearly be seen at a delay of +50 ms. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Fig. 3.

Illustration of algorithm decomposition into CUDA framework. (A) Output of the algorithm on a 2 min dataset recorded from human cortex. Each pixel corresponds to a covariance calculation of a window of data from the destination channel either before (top half) or after (bottom half) the source channel. This type of calculation can show nonstationary correlation measures, but takes an impractical amount of time on conventional CPUs. (B) Division of the algorithm into CUDA framework. Because each covariance calculation is atomic, we assign one calculation per CUDA thread. Each CUDA thread calculates one pixel in the top image. The number of lag calculations becomes the dimension of the block, up to 512 threads per block. If the number of lag calculations per time point is above 512, the work is divided into two or more kernel executions with block sizes designed to minimize “hanging”, no-calculation threads and blocks (white with red X). Because each time point shares a significant amount of the source and destination signal samples, both the source window and entire destination window required by the calculation is loaded into shared memory before execution. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Fig. 4.

Relative sizes and locations of the electrocorticography grids used in this study. Subject A was implanted with a subdural electrode array containing platinum disk electrodes 3 mm in diameter spaced 10 mm center-to-center, covering a total area of 490 mm². The implanted array covered primary hand motor cortex at the superior edge of the grid. A smaller grid was implanted in Subject B, with 1.5 mm electrodes spaced at 3 mm. In both subjects, orientation of the grid can be discerned by a blue dot in the corner containing electrode label 1. Cortical models for both patients were generated from pre-operative MRI scans. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Fig. 5. Experimental paradigm. The subject was presented with a blank silhouette of a hand. After 2 s, a finger would be highlighted (or pair, for pinch) and the subject would flex the indicated finger. After two seconds, the highlight would disappear and be replaced with the silhouette again.

Fig. 6.

Example results of covariance interactions in Subject B's smaller ECoG grid. (a) Index finger flexion vs. time. Higher values indicate more curl. (b) Plot of the interaction of a single electrode (red) with remote electrodes (green). (c) Covariance measurement plot, showing Time vs. Lag. Maximum finger flexion occurs at time $t = 0$. A period of high covariance is seen at the beginning of finger flexion at a lag of 98 ms. This signifies that after high gamma activity occurs in the source electrode (red), a covariant increase in high gamma is recorded in this electrode 98 ms later. Because the covariance is transient, it is specific and tightly coupled the finger flexion event. (d) This electrode shows period of coincident covariance during the onset of finger curl. At maximum finger flexion the coincident activity becomes shifted by 85 ms. (E) During finger curl, this electrode shows event-related covariance 53 ms before the source electrode (red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Fig. 7.

(A) In Subject A, unique interactions between sensory, motor and premotor can reliably be identified while areas not associated with volitional movement generation show no event-related covariance. (B) SMA areas (electrode 58) show faint but unique patterns of interaction with sensory areas (electrode 17) across digits. Motor (electrode 33) and sensory interactions show large event-related covariance across all digits. Areas that are expected to have no event-related covariance, such as parietal and lateral mouth sensory (1 and 28, respectively) show no event-related covariance for any digit.

Fig. 8. Spatially distributed activity patterns for Subject A. For source electrode 17, unique and event-related patterns of activity can be seen for each digit. In the case of pinch, covariance can be seen in remote electrodes that are a combination of thumb and index finger.

Fig. 9.

Illustration of spatially distributed patterns of activities for each finger across the microgrid implanted in Subject B. Source electrodes for covariance measurements are marked as red squares. Electrode 52 shows high event-related covariance with a large portion of the grid during thumb movements but few for other digits. Electrode 49 shows low but consistent activation across all digits, each uniquely spatially distributed. Electrodes 44, 28, 19 and 4 each show preferential event-related covariance for index, middle, ring and pinky fingers, respectively.

Table 1

STWC calculation on 100k random single-precision samples, with window size of 300 samples and pre-/post-lag of 300 samples.

	Mean time (s)	Std (s)	Relative performance to 1 Thread CPU
CPU (1 Thread)	23.13	0.3811	1.0×
CPU (2 Threads)	14.51	0.6448	1.59×
CPU (3 Threads)	13.23	0.1839	1.75×
CPU (4 Threads)	12.90	0.1304	1.79×
GPU	0.7337	1.667e-3	31.52×