

©Copyright 2014

Elliot Saba

A Fresh Look at Functional Connectivity Analysis

Elliot Saba

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Master of Science in Electrical Engineering

University of Washington

2014

Committee Members:

Les Atlas

Date

Adrian KC Lee

Date

Program Authorized to Offer Degree:
UW Electrical Engineering

University of Washington

Abstract

A Fresh Look at Functional Connectivity Analysis

Elliot Saba

Functional connectivity analysis attempts to detect methods of communication between groups of neurons in the brain that may or may not be directly physically connected. We briefly review current neuroscience methods for functional connectivity analysis, noting the similarities between the current state of the art as described in Miller et. al [4] and the recent work in the field of complex-valued statistics pioneered by Scharf et. al [7] and Picinbono et. al [5]. We apply the techniques of these complex-valued statistics to simulated and natural data, showing that these techniques are able to correctly and accurately detect a common mode of communication known as amplitude/phase coupling in the neuroscience community. This analysis is shown to reinforce the results of the study of in [4], with the added and substantial benefit of being completely automatic and not requiring human intervention to produce results.

Chapter 1

INTRODUCTION

1.1 Overview

The brain is the most mystifying and complex of all organs found in the human body. Understanding the mechanisms and processes within this bundle of neurons and grey matter that hosts the phenomenon known as consciousness has become a larger and larger focus of the biomedical community over the last century, especially with the advent of neuroimaging technologies. Techniques such as Functional Magneto-resonance imaging (fMRI) and Electro- or Magnetoencephalography (EEG/MEG) opened the doors to exploring the workings of a living brain by measuring the activity of specific areas in the brain over time. By recording and analyzing this activity during specific tasks, neuroscientists have been able to infer which areas of the brain are involved in the tasks. This analysis is known as functional connectivity analysis, as the areas of the brain that show correlated activity may not be physically connected but are nonetheless communicating in some fashion, collaboratively performing the task set before the subject.

Applications for methods to discover this sort of connectivity range from the purely scientific to medical and neuroengineering-oriented. Autism spectrum disorder (ASD) researchers hypothesize that functional connectivity could be a diagnosis method for determining whether a subject has ASD as well as a valuable tool for understanding the underlying cause for ASD in the first place [2]. Additionally, state of the art brain-computer interfaces used for direct neural control of external devices depend on an accurate understanding of what the signals picked up from the brain represent, a question functional connectivity analysis attacks directly by investigating the methods of communication within the brain itself.

Functional connectivity analysis methods and procedures remain an open area of research, including work on Functional Magnetic Resonance Imaging (fMRI) methods that infer spatially-separated yet correlated brain activity through blood-oxygen content, as well as cross-frequency coupling methods that operate directly on the electromagnetic fields created by neuron activity. In this thesis we will investigate cross-frequency coupling, categorized into several distinct models, each attempting to quantify interactions between two signals denoted as the discrete-time signals $x_1[n]$ and $x_2[n]$. Depending on the experimental and analytical paradigm being employed, these signals could either be derived from a single source signal $z[n]$, (as is the case when comparing two separate regions in frequency from a single spatial recording location) or could be derived from two separate source signals $z_1[n]$ and $z_2[n]$ (as is the case when comparing two separate spatial recording locations). When dealing with only a single source signal, we will denote it without any subscript, e.g. $z[n]$. The separate models as described below each represent a different mathematical relationship between the two signals, and each have different physiological reasons behind why they would be related by those mathematics. We do not attempt to explain the physiological processes behind the coupling, but instead content ourselves with developing methods to test each of these hypothesized models:

- *Amplitude/Amplitude coupling*: This represents the amplitude of one signal correlating with the amplitude of another. This is commonly referred to as amplitude modulation, or AM, and represents correlation between the amplitude envelopes of both $x_1[n]$ and $x_2[n]$.
- *Phase/Phase coupling*: This represents the phase of one signal correlating with the phase of another. This represents a linear relationship between the frequencies of $x_1[n]$ and $x_2[n]$, akin to harmonicity.
- *Amplitude/Phase coupling*: This represents the amplitude of one signal correlating with the phase of another. In this document we will use the convention that it is the phase of $x_1[n]$ and the amplitude of $x_2[n]$ that is coupled.




Figure 1.1: A simulated example of Amplitude-Phase coupling; this signal z contains a low-frequency component and a high frequency component, where the phase of the low-frequency component x_1 modulates the amplitude of the high-frequency component x_2 .

A simulated example of amplitude-phase coupling is shown in Figure 1.1, where a high-frequency oscillation is more likely to occur at zero sine phase of a low-frequency oscillation.

1.2 Our contribution

This document focuses on Amplitude/Phase coupling, constructing a novel estimator to detect this coupling, testing that estimator on synthetic data to ensure it captures the signal properties required of it, and finally testing it on natural data that is a priori known to have coupling embedded within it. We show that this analysis extracts the same information garnered by previous methods but most importantly without resorting to methods that require the intervention of a human which inevitably introduces unpredictable error.

Chapter 2

BACKGROUND THEORY AND OUR APPROACH**2.1 Modeling of the problem**

We model cross-frequency coupling as the interaction between two signals, $x_1[n]$ and $x_2[n]$. These signals can be disparate regions of frequency from a single source signal $z[n]$, or can be from two separate source signals $z_1[n]$ and $z_2[n]$, the mathematical analysis is the same. In this thesis we will focus on the case of the two signals originating from a single source $z[n]$ and representing disparate regions of frequency, although we reaffirm that the mathematics are identical across the two cases. The hypothesis we wish our estimator to test is that the phase of $x_1[n]$ is in some manner correlated with the amplitude of $x_2[n]$. We define our terms of “angle” and “amplitude” as follows:

- *Angle*: The complex angle of the analytic signal of $x_1[n]$:

$$\angle \hat{x}_1[n] \tag{2.1}$$

The analytic signal $\hat{x}_1[n]$ is constructed by passing $x_1[n]$ through the Hilbert transform, a transform that negates all negative frequencies from a signal, and then subtracting the result from the original signal:

$$\hat{x}_1[n] = x_1[n] + jH(x_1[n]) \tag{2.2}$$

Therefore, the resultant analytic signal $\hat{x}_1[n]$ contains only the positive frequencies of $x_1[n]$ and is complex-valued. Taking the angle therefore corresponds to the “phase” of the original signal, however we caution that this is only an accurate measure of instantaneous phase under extremely tight conditions [1]. We ensure that this measure is accurate in all simulations and analysis by constraining $x_1[n]$ to be narrowband by filtering it before applying Hilbert transforms.

- *Amplitude*: The absolute value of the analytic signal of $x_2[n]$:

$$\left| \hat{x}_2[n] \right| \quad (2.3)$$

This represents a time-varying spectral energetic measure, once again estimated through the Hilbert transform of $x_2[n]$. This spectral-energetic measure is widely used in signal processing for performing just this sort of analysis, however we point out that as we are taking an absolute value and hence forcing non-negativity, this measure does not represent amplitude; it is better expressed as magnitude or energy.

Using these two definitions, we can construct an intermediary signal $y[n]$, consisting of the phase of $x_1[n]$ and the magnitude of $x_2[n]$:

$$y[n] = \angle \hat{x}_1[n] \left| \hat{x}_2[n] \right| \quad (2.4)$$

This fits the polar form of a complex variable very naturally, as we have already segregated this quantity into magnitude and phase:

$$\begin{aligned} y[n] &= a[n] e^{j\theta[n]} \\ a[n] &= \left| \hat{x}_2[n] \right| \\ \theta[n] &= \angle \hat{x}_1[n] \end{aligned} \quad (2.5)$$

Therefore we will refer to the magnitude of $y[n]$ as $a[n]$ (or, equivalently, the magnitude of $x_2[n]$) and to the angle (or phase) of $y[n]$ as $\theta[n]$ (or, equivalently, the angle of $x_1[n]$). We will also refer to $a[n]$ and $\theta[n]$ as the “magnitude” and “phase” signals, respectively.

This composite signal $y[n]$ is constructed as described above primarily because it sidesteps a deep theoretical issue when attempting to detect phase/amplitude interactions; Rather than attempting to estimate the phase of $x_1[n]$ directly (an operation that is non-trivial at best, see [1] and then construct a measure to compare the estimated phase with the estimated amplitude of $x_2[n]$, this composite signal encodes both the phase and magnitude (a feature closely related to amplitude) of $x_1[n]$ and $x_2[n]$ respectively into a single signal. We are then able to correlate the magnitude and phase directly, using a field of mathematics known as circular statistics [5], which has enjoyed a recent rise in research focus over the last decade.




Figure 2.1: Left: A realization of a proper complex random variable, whose probability density function is rotationally invariant (“Circular”). Right: A realization of an improper complex random variable, whose probability density function is rotationally variant (“Non-circular”).

2.2 Circular Statistics Background

Circular statistics aim to quantify the relationship between a complex variable y and its conjugate y^* . This relationship is important because, as explained by Schreier and Scharf in [6], there exist significant differences between the class of complex-valued random variables that lack conjugate correlation (proper variables) and the class of signals that have some kind of correlation (improper variables). Schreier and Scharf motivate their findings by noting that improper complex random variables have second-order characteristics that are not completely described by their traditional covariance $E[yy^*]$ measure, rather they must be characterized by both their covariance and their complementary covariance, calculated as $E[yy]$. In this context, a lack of correlation between y and y^* translates intuitively to a random process with a probability density function in the complex plane that is rotationally invariant. Figure 2.1 shows a comparison between a proper and improper complex

random variable, and illustrates why proper random processes are sometimes referred to as “circular”.

Schreier and Scharf proposed a method to measure the “non-circularity” of a realization of a complex random process in [7] via a generalized likelihood ratio test. By measuring the complementary correlation $E[yy]$, and normalizing by the traditional correlation term $E[yy^*]$ of the complex-valued vector y , one can determine how improper a particular realization is. This measure, denoted ρ and defined as:

$$\rho = \frac{E[yy]}{E[yy^*]} \quad (2.6)$$

$$0 \leq |\rho| \leq 1 \quad (2.7)$$

has magnitude constrained to be within 0 and 1, where 0 represents a completely proper signal, and 1 represents a pathologically improper signal. This “non-circularity coefficient” serves as the basis for detecting coupling in our composite signal $y[n]$, creating a detector for interactions between the magnitude and phase of a complex random variable.

2.3 Our Method

Given the above background, we can merge the stated concepts and create a detector for amplitude-phase coupling. We list our algorithm as a series of steps, giving necessary implementation details yet leaving more in-depth analysis of results for later sections:

- *Separation of phase/amplitude portions*

First, the original signal(s) must be segregated into $x_1[n]$ and $x_2[n]$, which represent the signals from which we will derive the phase signal and magnitude signal, respectively. In this document, this is done using a Daubechies **db2** wavelet filterbank with 8 levels. We will take pairwise combinations of wavelet scales and use those pairs as our $a[n]$ and $\theta[n]$ signals, inspecting the interactions between each possible combination in our search for amplitude/phase coupling.

- *Construction of composite signal*

Using the segregated $a[n]$ and $\theta[n]$ signals, we combine them as shown in equation 2.4

to generate the composite signal y for each pairwise combination of signals chosen to be $a[n]$ and $\theta[n]$.

- *Calculation of non-circularity coefficient*

Using the Generalization Likelihood Ratio Test (GLRT) given in equation 2.6, for each composite signal we calculate the non-circularity coefficient. This therefore defines a “coupling matrix,” where each row and column correspond to a choice of wavelet filterbank bin for $x_1[n]$ and $x_2[n]$, respectively, and the intensity of each entry in the matrix corresponds to the “coupling coefficient” between those two choices of wavelet scales of $a[n]$ and $\theta[n]$.

We therefore have constructed an algorithm to generate a coupling matrix allowing easily automated analysis of signals to determine not only if there is coupling in a signal, but also between which regions in frequency the coupling occurs, as well as its strength.

Chapter 3

THEORETICAL/SIMULATED RESULTS

To verify the algorithm works as expected, we construct a pair of tests to determine the failure modes of the analysis. We show that the analysis is robust in the presence of noise, and that it does not generate false positives when analyzing signals without any inherent coupling.

3.1 *Creation of synthetic test signals*

To test the algorithm on signals without coupling, we generate z_{noise} as zero-mean, white Gaussian noise with standard deviation σ (Where a Gaussian random variable is hereafter denoted mathematically as $N(\mu, \sigma)$):

$$z_{\text{noise}} \sim N(0, \sigma)$$

Running the algorithm as outlined in section 2.2, this random process will be segregated by frequency, transformed into “phase” and “magnitude” portions, and then correlated via the GLRT to generate the coupling matrix W . Ideally, this signal should result in the zero matrix, as it should not detect any coupling in the test signal at all.

To test the algorithm on signals with explicit coupling, we generate z_{coupling} out of two signals: a low-frequency sinusoid with random amplitude and frequency modulation serving as the “phase” portion of z_{coupling} , and a high-frequency noise burst at times that line up with the phase of the low frequency portion. A realization of such a synthesized signal is shown in Figure 1.1, where the high-frequency noise bursts occur at sin phase 0 of the low-frequency component, and the noise burst is shaped via a Gaussian window. This signal is then added to noise, and passed through the wavelet filterbank described in 2.2 to separate the signal into low and high frequency components. We are then able to construct the coupling matrix from the pairwise combinations of low and high frequency components and plot its magnitude to detect regions in frequency that contain coupling.




Figure 3.1: Left: A coupling matrix calculated from z_{noise} , showing that no significant coupling is found in random noise. Right: A coupling matrix calculated from z_{coupling} , showing that significant coupling was found between the magnitude of high frequencies and the phase of a lower frequency.

3.2 Analysis of synthetic test signals

Coupling matrices from both of these test signals are shown in Figure 3.1. Note that the matrix calculated from z_{noise} is very low-energy, whereas the matrix calculated from z_{coupling} has significant energy in the top-right corner. This is consistent with the construction of the signals, as the top-right corner denotes a low frequency (high wavelet scale) for the phase signal, and a high frequency (low wavelet scale) for the magnitude signal.

3.2.1 Performance in noise

When analyzing Gaussian noise as a test signal that contains no coupling, it can be shown that the coupling matrix will converge to the zero matrix. This follows from tracing the input signal through the filtering operations until it is applied to the generalized likelihood ratio test, and showing the behavior of the test as a whole when applied to noise. Starting with an input of z_{noise} , which we assume grows eventually to infinite length, we filter it through a series of linear filters (wavelet filterbank, Hilbert transform), and then take the

absolute value or angle of the resultant signal:

$$a[n] = |h_1[n] * z_{\text{noise}}[n]| \quad (3.1)$$

$$\theta[n] = \angle(h_2[n] * z_{\text{noise}}[n]) \quad (3.2)$$

Where h_1 and h_2 represent the combined wavelet and Hilbert transform filters necessary to generate x_1 and x_2 , respectively, and the $*$ operator refers to signal convolution, rather than multiplication. As we know that z_{noise} is Gaussian noise, when it is Hilbert transformed its analytic signal will become a complex Gaussian with independent real and imaginary parts. Because a linearly filtered Gaussian random process is still a Gaussian random process, we can use the well-known result that taking the absolute value of the filtered quantity in 3.1 will result in a Rayleigh distributed random variable. Similarly, the angle of the filtered quantity in 3.2 will result in a uniformly distributed random variable.

Therefore, if we construct the composite signal $y[n]$ and perform the GLRT test on it, we can look at how this signal will perform in expectation when fed signals that conform to these random distributions:

$$y = a[n]e^{j\theta[n]} \quad (3.3)$$

Putting y into the GLRT, we see:

$$\rho = \frac{E[y[n]^T y[n]]}{y[n]^H y[n]} \quad (3.4)$$

$$= \frac{E[(a[n]e^{j\theta[n]})^T a[n]e^{j\theta[n]}}{E[(a[n]e^{j\theta[n]})^H a[n]e^{j\theta[n]}} \quad (3.5)$$

Because $a[n]$ is real, $a[n]^H = a[n]^T$. We also have the identity that $(e^{j\theta[n]})^H = (e^{-j\theta[n]})^T$. Using these, we can reduce the above to:

$$\frac{E[(a[n]e^{j\theta[n]})^T a[n]e^{j\theta[n]}}{E[a[n]^T a[n]]} \quad (3.6)$$

Therefore, we can immediately show that the denominator of the above fraction is a Rayleigh distribution squared, which becomes an exponential distribution, the mean of which is well-known. Assuming our initial Gaussian variable had standard deviation σ^2 , the Rayleigh distribution of $a[n]$ has parameter σ , and the exponential distribution therefore




Figure 3.2: The average energy of a W matrix analyzing noise converges to zero as the length of the noise signal grows to infinity. Error bars of ± 1 standard deviation are shown as dashed lines about the mean value, shown as the solid line

has parameter σ as well. Therefore, the expected value of the exponential distribution is simply σ , and we can replace the denominator of the above fraction with this value:

$$\frac{1}{\sigma} E[(a[n]e^{\theta[n]})^T a[n]e^{\theta[n]}] \quad (3.7)$$

Moving forward, we can re-express the expression within the expectation above as follows:

$$\frac{1}{\sigma} E\left[\sum_i a[i]^2 \cdot e^{2\theta[i]}\right] \quad (3.8)$$

Therefore, because $a[n]$ and $\theta[n]$ are by construction completely independent, we can separate them in the expectation:

$$\frac{1}{\sigma} \sum_i \left(E[a[i]^2] E[e^{2\theta[i]}] \right) \quad (3.9)$$

However, as $\theta[n]$ is uniformly distributed, $E[e^{2\theta[i]}] = 0$, and the whole expectation reduces to 0. This shows that feeding noise into this algorithm results in the zero matrix,

and as such this algorithm rejects noise as the signal extends to infinity. This is supported by experimental results such as those shown in Figure 3.2, where we can see the energy of the coupling matrix W converging to zero as the length of the noise being analyzed is increased.

3.2.2 *Addition of noise to coupled signal*

Combining these two experiments, we can analyze the behavior of the estimator acting upon a coupled signal in the presence of noise and see how the estimated matrix W deteriorates as the signal to noise ratio decreases. Simulations on this, (taking the synthetic signal as described in Section 3 and adding Gaussian white noise to it) are shown in Figure 3.3, where the power of the noise added to the signal is referenced to the power of the high frequency component coupled with the low frequency component. Although the signal is eventually lost due to the noise overwhelming the signal and therefore causing the phase and amplitude signals even after the filtering to lose all coherence with each other.

3.2.3 *Analysis of phase*

Finding the presence of coupling is not the only aim of cross-frequency coupling however, as it is also desirable to characterize the phase of the low frequency for which the high frequency component achieves large magnitudes, the so-called “preferred phase”. Attentive readers may expect that the angle of ρ from Equation 2.6 could be used for this task, however it is important to remember that ρ is directly proportional to the expected value of y^2 , therefore if we assume the form $y[n] = a[n]e^{j\phi[n]}$, we can see that the phase of the result will be doubled, and it is therefore impossible to uniquely determine the original phase of the distribution from simply the phase of ρ .

Alternative methods exist of course, and as it is not the purpose of this thesis to invent new methods of estimating the average phase of circular data, we will refer to other predefined methods such as usage of the Fisher-Von Mises distribution [3] or angular histograms, where the data is grouped into bins of a certain width in phase.

We show a scatter plot of the absolute value of the simulated composite signal $y[n]$




Figure 3.3: Left: Coupling matrix of synthetic signal with high-frequency signal twice as powerful as the background Gaussian noise. Right: Coupling matrix of synthetic signal with high-frequency signal at the same power the background Gaussian noise.




Figure 3.4: Plotting the magnitude of the composite signal against the phase of the composite signal, we recover a compelling visual of the nonuniformity (and hence noncircularity) of the composite signal $y[n]$

versus the phase of that same signal in Figure 3.4, as well as a histogrammed version of the same plot so as to more directly compare with previous work in which such an operation is typical. In both cases we can see the readily apparent nonuniformity of the distribution of magnitude across phase, thereby confirming some kind of magnitude-phase coupling.

Chapter 4

EMPIRICAL RESULTS

In this section we analyze naturalistic data taken from experimental subjects via electrocorticography which has been previously published by Miller et. al and is known to contain cross-frequency coupling in the form of amplitude/phase coupling [4]. We compare the previous paper’s methods to our own, and show similar results.

4.1 Data and Previous methods

In Miller et. al [4], electrocorticographic data was collected from epileptic patients while performing a visual search task. Data was recorded from the occipital lobe via multiple electrodes, sampled at 1000Hz. The data from the electrodes was demeaned (across electrodes, e.g. the average value of all electrodes was subtracted from each electrode), and filtered using a wavelet filterbank into separate frequency regions. Low frequency regions were manually inspected for peaks and valleys and annotated by phase. Miller et. al then binned high-frequency log-magnitude envelopes (as calculated by the absolute value of the analytic signal of the high-frequency region being analyzed) by the binned annotated phase of the low-frequency region.

This analysis resulted in “coupling palettes”, showing uniform distributions of magnitude across phase when no coupling existed, but a nonuniform distribution when some form of amplitude/phase coupling existed, similar to the histograms shown in Figure 3.4, but plotted across multiple choices of magnitude signal frequency. Compared with our proposed method, there are many similarities, however we aim to contribute to this field in two main areas:

First, our method is completely automatic and does not rely on any phase or time estimation on the part of the algorithm or humans, when determining the phase estimate, or low-frequency oscillations in the analyzed signal. The algorithm does not attempt to estimate instantaneous phase, either automatically or manually, and as such does not suf-




Figure 4.1: A coupling matrix calculated from naturalistic data, showing the amplitude/phase coupling inherent in the signal

fer from the problems inherent in estimating that quantity and does not require human intervention.

Second, our method, apart from initial wavelet filterbank choice, is quite nonparametric. Binning does not occur after the signal has been segregated, and specifically the detection of coupling (in our algorithm, done by the generalized likelihood ratio test) does not depend on an arbitrarily-chosen parameter such as number of bins in time/phase.

We note that the discussion in section 3.2.3 regarding analysis of phase is independent of the detection of coupling in the first place, and therefore should not be regarded as a contradiction to the above statements. Section 3.2.3 aims to show that analysis of coupling data from our method contains the same phase preference information as in previous methods, however the steps needed to gain that data are simpler and better theoretically justified in our case due to the lack of parameters and human intervention.




Figure 4.2: Plotting the magnitude of the composite signal against the phase of the composite signal, we recover a compelling visual of the nonuniformity (and hence noncircularity) of the composite signal $y[n]$

4.2 Application to natural data

When run on real data from [4], analysis was performed on data as described in 2.2, and an example W matrix is shown in Figure 4.1. This result shows that coupling between the two frequency regions as reported in [4] was successfully found, however Miller et. al furthermore showed the phase at which the coupling was found between these two frequencies. A method to calculate this (very similar to that used in Miller et. al) is to bin the magnitudes of the data by their phase, and this very approach is shown in Figure 4.2. Unfortunately, Miller et. al do not provide sufficiently fine-grained results for me to check this particular result against their previously published results, however we are confident that given the theoretical experiments from Section 3.2.3 this apparent coupling is not the product of random chance.

Chapter 5

CONCLUSION

In conclusion, we have proposed a novel method for detecting and analyzing cross-frequency coupling between two frequencies in signals, and have shown that it provides satisfactory results on both simulated and naturalistic data. We have compared it to a previous algorithm, verifying its output and detection capabilities, while still retaining some advantages over the previous method such as a reduction in parameterization, a theoretically justified resilience against noise and most importantly a move away from the necessity of human intervention in the analysis process. These advantages open up the possibility of mass-analyzing data in a way that is not possible when human effort is necessary to complete the analysis and can possibly open doors in neuroimaging analysis that have previously been closed, where the amount of data to be analyzed was previously simply too high.

In future work, it would be beneficial to investigate more principled methods of determining optimal frequency segregations of source signals. Currently, we use a wavelet filterbank, and although this works well there is currently no satisfactory argument for why wavelets should work better than other, alternative methods of carving up the frequency domain of neuroimaging signals. We also encourage deeper investigation of cross-frequency coupling when the source signals do not originate from the same physical location in the brain. In this thesis, so as to be able to compare our results with previous work we have limited ourselves to signals originating from the same location, as that is the extent of the analysis previously performed on this data, however the end goal of cross-frequency coupling, and the most intriguing applications are concerned with the communication of disparate frequencies from disparate locations. The application of the algorithms and mathematics contained within this thesis follows naturally from the single-location case, and as such this analysis would likely yield interesting scientific results concerning the nature of our brains.

ACKNOWLEDGMENTS

To Prof. Les Atlas, for his continual encouragement, support and challenges to take my thinking to a higher level.

To Prof. Adrian KC Lee, for always knowing what I needed to do better than I did myself, and for having the grace to let me discover it anyway.

To Dr. Pascal Clark, for treating every blunder of mine as an opportunity to teach me a better way.

BIBLIOGRAPHY

- [1] B. Boashash. Estimating and interpreting the instantaneous frequency of a signal. pages 520–538, 1992.
- [2] S. Deb and B. Thompson. Neuroimaging in autism. *British Journal of Psychology*, 173:299–302, 1998.
- [3] R. Fisher. Dispersion on a sphere. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 217(1130):295–305, 1953.
- [4] Kai J Miller, Dora Hermes, Christopher J Honey, Mohit Sharma, Rajesh P N Rao, Marcel den Nijs, Eberhard E Fetz, Terrence J Sejnowski, Adam O Hebb, Jeffrey G Ojemann, Scott Makeig, and Eric C Leuthardt. Dynamic modulation of local population activity by rhythm phase in human occipital cortex during a visual search task. *Frontiers in human neuroscience*, 4(October):197, January 2010.
- [5] B. Picinbono. On circularity. *Signal Processing, IEEE Transactions on*, 42(12):3473–3482, 1994.
- [6] P.J. Schreier and L.L. Scharf. Second-order analysis of improper complex random vectors and processes. *Signal Processing, IEEE Transactions on*, 51(3):714–725, 2003.
- [7] P.J. Schreier, L.L. Scharf, and Alfred Hanssen. A generalized likelihood ratio test for impropriety of complex signals. *Signal Processing Letters, IEEE*, 13(7):433–436, 2006.