Analyzing Neural Time Series Data: Theory and Practice

Mike X Cohen

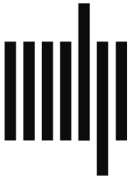

The MIT Press

From The MIT Press

MITCogNet

This book was set in ITC Stone Serif Std by Toppan Best-set Premedia Limited, Hong Kong. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Cohen, Mike X, 1979–
Analyzing neural time series data : theory and practice / Mike X Cohen.
pages cm.—(Issues in clinical and cognitive neuropsychology)
Includes bibliographical references and index.
ISBN 978-0-262-01987-3 (hardcover : alk. paper) 1. Neural networks (Neurobiology) 2. Neural networks (Computer science) 3. Computational neuroscience. 4. Artificial intelligence—Biological applications. I. Title.
QP363.3.C633 2014
612.8'2339—dc23
2013016461

10 9 8 7 6 5 4 3 2 1

37 Recurring Themes in This Book and Some Personal Advice

There are several themes in this book that permeate the text, figures, and analyses. Some of these themes were repeatedly made explicit; others were implicitly suggested. This chapter reiterates and expands on these themes. I also give some personal advice about analyzing cognitive electrophysiology data. The advice is based on my personal experiences, and you are free to disagree.

37.1 Theme: Myriad Possible Analyses

There seems to be a limitless parameter/analysis space within which to analyze data. This can be intimidating to newcomers in the field because they do not know which analyses should be used in which circumstances. And this may be fodder to skeptics unfamiliar with the analyses, who will claim that any finding desired can be obtained by "torturing the data" enough (although I doubt this is true).

If you are unsure which analyses to perform, follow the analysis protocol of a published paper on a similar topic. If there are no relevant publications, or if you think the published analysis protocols are not optimal for your study, start with the simple analyses and apply more complex analyses only when necessary.

Analyses are like tools in a toolbox: you should pick the right one(s) for the job. Even your favorite method might not always be appropriate, much as a sledgehammer might not make the best fly swatter. Whether an analysis is appropriate depends almost entirely on the research question. Data analyses should be as sophisticated as they need to be to answer the question, and you should avoid using complex methods when simpler methods will answer the question equally well. If your preference is to pick a method and then find a problem to which to apply that method, you should try to find an appropriate research question for which that method is particularly well suited.

37.2 Advice: Avoid the Paralysis of Analysis

Analyzing data is fun and satisfying (for some people). It provides both instant and long-term gratification and gives a sense of achievement and closure when you solve a data analysis problem. But for many cognitive electrophysiologists, particularly those toward the "cognitive" end of the spectrum of figure 1.1, analyzing data should be the means to the end, not the end in itself. The purpose of analyzing EEG data should be to extract meaningful patterns that are embedded in the EEG signal but are difficult to observe in the raw signal, and the result of data analyses should be empirical support for or against a statement concerning some aspect of brain function.

The potential problem with data analyses is that the search space of electrophysiology data is so huge that there are limitless possibilities for new or different ways to analyze the same data. Thus, there is a risk of ceaselessly analyzing and reanalyzing your data without getting any new information from the data. If your head is in the data-analysis clouds, try to keep your feet on the purpose-of-the-study ground.

On the one hand, it is good to be cautious and check your analyses, and it is good to confirm a result using alternative analysis methods. But on the other hand, at some point new analyses will take more of your time and resources without giving you anything back in return. At some point, you need to accept what you've done with that dataset and move on.

37.3 Theme: You Don't Have to Program Your Own Analyses, but You Should Know How Analyses Work

Programming your own analyses increases your freedom and flexibility. It allows you to have more control over what analyses you can perform, and it allows you to custom-tailor analyses to the idiosyncrasies of your experiment design and your data, and to select the most appropriate analysis parameters for your hypotheses. It also helps you interpret the results more appropriately because you understand at a mathematical level what happened to the data.

There is nothing wrong with using existing and user-friendly analysis packages or running Matlab scripts that someone else wrote. It is important to understand at a conceptual and mathematical level what happens to your data between pressing the "record" button during data collection and making publication-quality figures. Without an understanding of at least the basics of how the analyses work, there is a danger of using inappropriate analysis parameters, misinterpreting results, or doing something wrong without even knowing it.

You should try to understand at least conceptually what happens to your data, the advantages and disadvantages of each analysis method, the key parameters and what effects those parameters have on the results, when it is appropriate and inappropriate to use that analysis, and how the results can be interpreted. This level of understanding is not the same as—and is more important than—being able to write out all the equations on a blank piece of paper or being able to write Matlab code from scratch to do the analyses.

37.4 Advice: If It Feels Wrong, It Probably Is

Another major benefit of understanding how analyses work is that you develop an intuition for when a result doesn't "feel right." This intuition builds with experience as well as mathematical knowledge, but by understanding the mathematical bases of analyses and how they are implemented, you will develop a kind of "sixth sense" about when something went wrong or when there are artifacts in the data. This is no metaphysical awareness but rather a set of mental templates for what analysis results should look like. Furthermore, when something goes awry, by understanding how analyses work, you will have some good guesses about what the problem might be and how to solve it.

37.5 Advice: When in Doubt, Plot It Out

A picture is worth a thousand words. If you are not sure what a particular Matlab function does, plot the data before and after applying that function to the data. This will help you develop an intuitive feeling of what the function does to the data. The Matlab commands `plot`, `hist`, and `imagesc` are the most useful functions for understanding what happens to your data. Furthermore, if you suspect that something is wrong with your results or with the analyses, plot the data at each step of the analysis until you see a problem. NaNs and Infs will produce empty spaces in figures. Plotting and looking at data are the best ways to understand how Matlab functions work, how your data are transformed, what mathematical equations mean, and what may have gone wrong with the data or with the analyses.

One advantage of software programs such as Excel and SPSS is that the data are all on the screen in front of you. You cannot avoid seeing the raw data. Thus, a disadvantage of software programs like Matlab and R is that the data are hidden, in that you won't see them unless you make a specific effort to look at them. This distinction between being visually confronted with your data and having to make an effort to see the data is akin to a distinction between bottom-up and top-down processing. If there is an outlier in, for example, reaction time data, you are likely to notice this outlier via bottom-up visual capture if your

data are in Excel, but you will need to exert top-down control to notice this outlier in Matlab (e.g., by plotting the reaction times of all trials). Most or perhaps all EEG analysis packages have visualization routines that allow you to look at and scroll through your raw data as well as results of the analyses. The importance of looking at raw and processed data cannot be understated, both in terms of making sure the data are clean and in terms of understanding the final results. Thus, try to employ top-down data inspection strategies because the bottom-up processes might be insufficient.

37.6 Advice: Know These Three Formulas like the Back of Your Hand

There are three equations that form the backbone of the majority of EEG time-frequency-based analyses. If you memorize these equations, you will be able to understand most time-frequency analyses. You should be so familiar with these equations that you can recite them at any time of day, in any mood you happen to be in, or in nearly any state of consciousness.

Sine wave: $A\sin(2\pi ft+\theta)$

Euler's formula: $Me^{ik} = M[\cos(k)+i\sin(k)]$

Gaussian: $e^{-t^2/2s^2}$

If you become very familiar with these equations, you will be able to see them embedded in seemingly complex equations that otherwise might be too difficult to understand. Here are a few examples: the discrete time Fourier transform is a combination of a sine wave and Euler's formula; a complex Morlet wavelet is a combination of all three formulas; wITPC and phase-amplitude coupling are Euler's formula. Obviously, these are not the only three equations that govern all EEG time-frequency-based analyses, but these are the main ones.

37.7 Theme: Connectivity over Trials or over Time

Many connectivity analyses can be done over time or over trials. Although the math is often very similar, the interpretations and the kinds of results you might obtain can be different. Connectivity over time increases sensitivity to simultaneous coupling and is better for high-frequency connectivity that might not be phase-locked to an experiment event. On the other hand, it has increased susceptibility to volume conduction, has relatively poor temporal precision, and may reflect tonic connectivity in addition to task-related connectivity. Connectivity over trials provides a better link to the timing of task events and has higher temporal precision. On the other hand, it can be used only with task data, is more affected by

having a small number of trials, and is less sensitive to detecting high-frequency connectivity that might not be phase-locked to the time = 0 event. Unfortunately, many publications do not specify whether their measure of connectivity was performed over time or over trials. Try to be clear about this when describing your analyses.

37.8 Theme: Most Analysis Parameters Introduce Bias

Nearly all analysis parameters introduce some kind of systematic bias into the results. Some of these biases are related to the trade-off between frequency and temporal precisions, and some biases are related to the kinds of findings you will observe or be blind to (e.g., phase-lag-based vs. phase-clustering-based connectivity measures or phase-locked vs. non-phase-locked activity). How concerned you need to be about these biases depends on the kind of bias and on the effects you are testing. In some cases, in comparing results across conditions, electrodes, time segments, and so on, the effects of biases are diminished or eliminated because the biases affect all results and are therefore minimized during subtraction or statistical comparison. This is why the same set of parameters should be applied to all data. For resting-state data or task data in which only one condition is examined, you should be more careful about interpreting the results in light of any biases that may have been introduced by analysis parameters. If you are using a new analysis or are unsure which parameter settings to use, you could perform the analysis several times using different parameter settings. If you do this, report in the publication what ranges of parameters you tried and whether that affected the patterns of results. This will not only help readers evaluate your results, it will also help other researchers, who are in a similar position as you are, to know which analysis parameters to use.

37.9 Theme: Write a Clear Methods Section so Others Can Replicate Your Analyses

This should be an obvious point, but it is worth stressing: write your Methods section clearly so that other scientists will be able to replicate your analyses. That does not necessarily mean that you need to write every minor detail, such as the formula for a Fourier transform. But there should be enough detail that someone with a reasonable amount of knowledge about analyzing EEG data will be able to reproduce your analysis protocol. If you use toolboxes, state which functions or methods were used and note any nondefault parameters. If you wrote your own code for the analyses, consider making the code publically available on a website, as an appendix or supplemental online section in the publication, or as an extension to an analysis toolbox such as eeglab or fieldtrip. If the analysis protocol was complex,

write a step-by-step list of what was done. When in doubt, have a colleague read the Methods section and decide whether he or she feels the method could be replicated. It is better to include too much than too little information, particularly for journals that either do not have strict word limits or that allow supplemental online material. If the method is novel, difficult, or involves many steps, consider making a figure to illustrate the procedure for how the analysis was conducted. There are recently published suggestions for what details to report in Methods sections of studies using MEG (Gross et al. 2013). That paper was written specifically with MEG studies in mind, but some of the suggestions are also applicable to EEG and LFP studies.

Be open and honest about the limitations of your methods and analyses that are relevant to the results. No method is perfect, and limitations are nearly always balanced with advantages. For example, the surface Laplacian increases spatial precision and allows connectivity results to be better interpreted at the electrode level, although this comes at the expense of filtering out some low-spatial-frequency features of the data. Not all limitations of a method are relevant to your study. For example, if your study concerns parietal alpha oscillations and attention, the spatial precision of EEG might not be a relevant limitation. However, if your experiment design is modeled after an fMRI study that suggests an important role for the thalamus, the spatial precision of EEG might relevant to mention. Another example is that time-frequency decomposition methods decrease temporal precision, but this may not be relevant unless your hypotheses involve testing for minor differences in timing between conditions.

37.10 Theme: Use Descriptive and Appropriate Analysis Terms

I thought carefully about the analysis terms used in this book and decided to use terms that are brief descriptions of the mathematical aspects of the method, rather than terms that are interpretations of the results. Some terms were avoided because they are ambiguous (for example, “spectral perturbation” can mean a change in power, phase, connectivity, band-specific graph properties, or other spectral features).

One could argue that it doesn’t matter what the method is called as long as the formulas are written and the procedure is clearly described. This rose-by-any-other-name argument is understandable, and it certainly is the case that the results and what they mean for brain function are important, not the short name given to a set of equations. But from a practical standpoint, using different terms in different papers to refer to the same analysis hinders efficient evaluation of results and also hinders cross-study comparisons. Ideally, when someone says or writes the name of a method, you should know exactly what method they

mean. In other words, there should ideally be a one-to-one mapping between the name of a method and the math behind that method. This is often the case in other branches of statistics. Consider the following terms: correlation, ANOVA, factor analysis, structural equation model. There is no ambiguity about which sets of equations are employed when someone says that he performed an ANOVA. Cognitive electrophysiology should be this exact, but it is not. If someone says that there was an increase in alpha synchronization, you don't know whether that person means an increase in power or phase-based connectivity, and these two analyses have very different interpretations, putative neurophysiological origins, theoretical implications, and methodological concerns.

The point is not to come up with perfect terms; at the end of the day, terms are merely words, and their connotations are influenced by semantic, cultural, linguistic features and personal preferences. I think Lewis Carroll best expressed the idea that there is an arbitrary relationship between words and their real-world referents (specifically through the character Humpty Dumpty in *Through the Looking Glass*). My point here is not that we should struggle to find the most linguistically perfect expressions for sets of mathematical equations that are applied to data, but rather, that there should be widespread agreement that a particular term implies a particular analysis. For example, you might think that the term "receiver operating characteristic" is a poor description of that analysis, but when you use that term (or its abbreviation ROC) to an audience with a statistics, psychology, or engineering background, it is very likely that they will know to which analysis you refer.

In my opinion—and others may disagree—the terms used to describe analyses should refer to the mathematical procedures that underlie those analyses, not to an interpretation of what the results of that analysis might mean in a neurophysiological sense. Interpretations of results are subjective, incite collegial disagreement, change over time as theories develop and new data are acquired, and can be incorrect or misguided. Mathematical procedures, on the other hand, are simply equations and data transformations and will remain the same regardless of what contemporary scientists believe the results of those analyses imply about brain function.

37.11 Advice: Interpret Null Results Cautiously

Null results can be difficult to interpret in any branch of science. There are many reasons why you might obtain a null result, including poor experiment design, low-quality data, too few trials, inappropriate analyses, and others. In this sense, null results from EEG analyses can arise for the same reasons that null results can arise in, for example, a psychology questionnaire study.

However, there are additional considerations for null results with EEG data, in part because EEG measures only a fraction of all brain activity and in part because, despite measuring only a fraction of brain activity, the EEG signal is multidimensional. You can obtain a null result simply because you looked at the wrong feature of the data, even if the experiment, data collection, and analyses were properly conducted.

For example, if a time-frequency plot of baseline-corrected power is all green (values near zero) in the theta band, it is inappropriate to conclude, "There was no theta activity in the brain." There could have been neurons firing in the theta band, there could have been theta activity that was spatially asynchronous so as not to generate an electrical field large enough to be measured by the electrode, there could have been deep theta generators that were filtered out during surface Laplacian, there could have been task-related theta activity in phase but not in power, and so on.

The appropriate interpretation of null results should be close to the analysis and should be anchored with a significant result if possible, which suggests that there was sufficient statistical power to detect a result if it were present. For example, "Theta power was not different between conditions A and B, whereas it was significantly greater in condition B compared to condition C."

It is also important to consider the situations in which a null effect would be observed with each particular analysis. For example, an absence of a change in connectivity as measured by phase-lag-based measures can indicate that there was no connectivity or that the connectivity had a near-zero- or -π phase lag. If an ERP is nonsignificantly different between two conditions, this may indicate that the same neurocognitive process was employed in those two conditions, or it could be that the neurocognitive processes that differed between the two conditions was non-phase-locked and therefore not visible in the ERP.

You can be more confident about interpreting null results if you can demonstrate (1) that the data are of sufficient quality and that there is sufficient statistical power to detect a result if it were present, and (2) that the analyses applied are appropriate to address the research question. Finally, if the interpretation of the null result is important for interpreting other results, show that null result in a figure in addition to reporting statistical *p*-values. That way, readers can judge whether the null result is going in a particular direction and could be significant if there were more or cleaner data, or whether the result is a true null effect and unlikely to be significant even if more and higher-signal-to-noise data were obtained.

37.12 Advice: Try Simulations but Also Trust Real Data

The obvious advantage of simulations is that there is a ground truth that you define, and you know whether the analysis worked because you can compare the results to the simulated

data. Simulations are also useful to illustrate the effects of extreme data on analyses because extreme data might be difficult to find in real data.

However, simulated data often do not look or behave like real data, particularly if the simulations do not contain realistic amounts of noise with realistic time-frequency characteristics. Therefore, if you are trying to understand how a method works, or if you are extending existing methods, you should use both simulated data and real data. If possible, use data from an often-used task that produces replicable results so this pattern can be considered an empirical ground truth.

37.13 Advice: Trust Replications

No dataset is perfect, and all datasets have idiosyncrasies and quirks. These may be due to subject sampling error or to experiment design features. One can make a distinction between a statistical false alarm that arose because of, for example, a lenient statistical threshold, and a result that is true in one sample but may not generalize to another sample. Replications are important because they demonstrate the generalizability of an effect beyond what evidence can be obtained from a statistical measure in one dataset. Of course, this does not mean that findings reported in only one study should be ignored or not trusted. Rather, findings should be taken more seriously after they have been replicated in several different experiments, ideally using different experiment designs and analysis approaches, and by different research groups.

This is another reason why using clear analysis terms and publishing tables of activation will facilitate identifying replicable effects across different research groups and experiments.

37.14 Theme: Analyses Are Not Right or Wrong; They Are Appropriate or Inappropriate

This theme was made explicit in most chapters, with such wording as *should you <insert question>? It depends on the goal of your study*. . . . If you have nothing better to do on a Saturday night, you can probably think of some data analyses that are completely wrong in all situations. But within the realm of reasonable analyses that a scientist would consider applying to his or her data, whether a particular analysis method should or should not be used depends on the goals of the study, the experiment design, the quality of the data, and the other analyses planned. Analyses that are appropriate in some situations might be inappropriate in other situations. For example, if you have a hypothesis about condition differences in phase-based connectivity between one frontal electrode and one parietal electrode, ISPC is a more appropriate measure than phase-lag-based measures because ISPC is more sensitive to detecting connectivity and condition differences, and the result can be inspected for possible

contamination of volume conduction. On the other hand, if you are examining whether individuals with autism have increases in short-range connectivity compared to matched controls, phase-lag-based measures are more appropriate than ISPC because some of the ISPC results may be contaminated by volume conduction, and it is impractical to examine each of hundreds or thousands of connectivity results for potential artifacts.

This is also why clear writing in the Methods section is important. It is possible that your analysis protocol is justifiable and appropriate, but because the description of your analysis protocol was not clear enough, reviewers and readers suspect that it was not done appropriately or, worse, that you intentionally biased the data selection procedure to obtain a specific result. For example, consider the following passage of an imaginary Methods section:

> In the first step of the analysis, we averaged time-frequency power from all conditions and selected a time-frequency window in the gamma band, based on a statistical comparison between time-frequency power over all conditions versus the prestimulus baseline. Note that because this selection procedure was based on the condition average, it is orthogonal to any possible condition differences. In the second step, we extracted power from this time-frequency region separately for each condition and each subject. These values were then entered into a 2 (visibility: low vs. high) × 3 (feedback condition: punishment, neutral, reward) ANOVA.

Now consider how this same procedure could lead readers to become suspicious because of unclear writing: "Condition differences in gamma power were tested where there were significant gamma power effects."

There are useful guidelines for knowing when certain methods are appropriate or inappropriate, which were (I hope) made clear in this book when each method was introduced and discussed. Ultimately, however, it is up to you to make that decision and justify it. If you are unsure whether a particular method is appropriate in your situation, ask colleagues or search for how that method was used in peer-reviewed publications. Keep in mind, however, that the published application of a particular method does not necessarily mean it was appropriately applied.

37.15 Advice: Hypothesis Testing Is Good/Bad, and So Is Data-Driven Exploration

Strict hypothesis testing is good: it is theory driven, has implications for theories and models, will likely help you design better experiments, and increases statistical sensitivity because it minimizes multiple-comparisons problems. Hypothesis-driven analyses also tend not to be very time-consuming because only a small number of analyses need to be performed. On the other hand, strict hypothesis testing is constraining and analogous to horses that wear blinders so they don't get distracted by their peripheral vision. Perhaps the most interesting

and most insightful results in the data are not the ones predicted by the limiting and probably incorrect theory (all theories are wrong . . .).

Data-driven exploration is also good: it offers freedom to discover patterns of results, provides an avenue for the more number-crunching-oriented scientists to expand and develop new methods, and it facilitates new discoveries unconstrained by the blinders of theories. Data-driven exploratory studies tend to be time-consuming, in part because there is so much multidimensional data to inspect and in part because many analyses are often performed as the researcher thinks "perhaps the best result is still hidden and will be revealed by a fancier analysis." On the other hand, data-driven exploration can be like fishing in a polluted lake (you might catch a fish or you might catch a plastic bag), requires appropriate correction for multiple comparisons that might push the most interesting and important results below the statistical threshold, and may provide results that are uninterpretable and are consigned to the basement of science because the finding is not relevant to any theory, model, or other research finding.

There is no correct position to take on the issue of hypothesis testing versus data exploration, and there is no optimal balance between them. The appropriate levels of hypothesis testing and data exploration should be based on your preferences, data analysis skills, patience, and orientation toward theory versus data. They are also based on the field in which you are working, including how much research has already been done in that area, what previous findings have shown, and whether there are theories detailed enough to make relevant predictions that can be confirmed or disconfirmed. In many cases simply too little is known about the electrophysiological dynamics of brain function and its relation to behavior, brain structure, and the body to perform only hypothesis-testing analyses.

These two approaches are not mutually incompatible. You can design an experiment to test specific hypotheses and then perform additional data-driven analyses to explore what else might be lurking in the data, waiting to be discovered. The other way around is possible as well, though less straightforward. If you design an experiment for data exploration, it is inappropriate to test specific hypotheses on the data after seeing the results, but it is appropriate (and good science) to design a new experiment to confirm the findings you observed in the exploratory analyses.

37.16 Advice: Find Something That Drives You and Study It

Science is hard. It requires time, energy, patience, perseverance, and self-discipline. For many scientists, science is not just a job, not just something you do during the day to earn money. Instead, it is a career, a passion, and a lifestyle. Whether you love or hate your life as a

scientist will affect your happiness and life satisfaction. Pick a research topic that fascinates you and a method that draws you. And keep long-term goals in mind. You won't love every single minute of being a scientist, but very often, the parts you don't like will help get you to the parts you do like. This advice is not just about preserving your sanity: it is also about becoming a better scientist. If you are uninterested in your research topic, you will probably put in the minimum required effort to turn it into a publication or conference poster presentation that others may not read past the abstract. If you are fascinated by your research topic, you will be motivated to do high-quality, creative, and progressive scientific research.

If you need some inspiration for scientists persevering—and ultimately triumphing—in the face of repeated rejections, setbacks, and constraints, consider reading the first chapter of *Roving Mars: Spirit, Opportunity, and the Exploration of the Red Planet*, by Steve Squyres (2005), who was the principal investigator of the mission to put mechanical rovers on the planet Mars.

37.17 Cognitive Electrophysiology: The Art of Finding Anthills on Mountains

Perhaps in an ideal world (at least, the cognitive electrophysiologist's ideal world), this book would not be necessary. Data analysis toolboxes and software programs would also not be necessary. You could simply put a few electrodes on one subject's head, test a few trials per condition, look at the raw EEG data with no processing or analyses, and understand the complex spatial-temporal-spectral landscape of cortical electrophysiological dynamics. There would be no noise, confounds, or alternative explanations to rule out, and there would be no sophisticated analyses, statistics, or probabilities to compute and base speculative inferences on. You would just look at the data and understand how the system works.

Regardless of whether your ideal world includes performing EEG research and looking at data, this is not the world we have. EEG signals are noisy, findings are sometimes hidden in dimensions that are difficult to visualize without filters or other data transformations, results can be infected with artifacts or may be statistical false alarms, and interpretations are usually based on probabilities, inferences, and speculations.

Good science—in addition to being based on theory, previous research, careful experiment design, and appropriate data analyses and statistics—is also about discovering and appreciating subtlety. And the appreciation of subtlety becomes more important with time as basic findings become established fact. This is good; it reflects progress. For example, you can no longer publish a paper in which you demonstrate that the human occipital lobe is involved in vision. Instead, you have to find more subtle and interesting features of the computations performed in visual processing areas and the neural implementations of those

computations. Often, the biggest and most obvious results are not the most interesting and insightful results. That doesn't mean you should ignore the obvious results—they are often important for data quality demonstrations and for characterizing the overall landscape of cortical electrical dynamics. But in many situations the big obvious effects are not the ones that provide novel insights into neurocognitive function.

Discovering a subtlety in the brain seems as though it should be easy, but it is not. Subtleties can be hidden in noise or hidden under a much more robust but less interesting result. You might have almost discovered a subtlety but had the wrong condition comparison, frequency band, or time window. Discovering subtleties involves considering theories, previous research, and openness to letting the data "speak for themselves." Science, or at least cognitive electrophysiology, can sometimes be more of an art than a science.