



Charting a future for fMRI in communication science

Benjamin O. Turner, Richard Huskey & René Weber

To cite this article: Benjamin O. Turner, Richard Huskey & René Weber (2018): Charting a future for fMRI in communication science, Communication Methods and Measures, DOI: [10.1080/19312458.2018.1520823](https://doi.org/10.1080/19312458.2018.1520823)

To link to this article: <https://doi.org/10.1080/19312458.2018.1520823>



Published online: 08 Oct 2018.



Submit your article to this journal [↗](#)



View Crossmark data [↗](#)



Charting a future for fMRI in communication science

Benjamin O. Turner^a, Richard Huskey^b, and René Weber^c

^aWee Kim Wee School of Communication and Information, Nanyang Technological University, Singapore, Singapore;

^bSchool of Communication, The Ohio State University, Columbus, OH, USA; ^cDepartment of Communication, UC Santa Barbara, Santa Barbara, CA, USA

ABSTRACT

Neuroscientific investigations into communication phenomena using functional magnetic resonance imaging (fMRI) are becoming increasingly popular in communication science. This presents opportunities for new discoveries, but also for the rapid spread of questionable practices. Here, we look to the future of fMRI in communication science: first, highlighting and advocating for several relatively new methods that should enable communication scholars to address novel research questions; and second, pointing out various controversies or pitfalls that exist in the use of several of the more widely used fMRI analysis methods within the field. Given the rapidly changing nature of the fMRI analysis landscape, such reflection is an essential part of being a good scholar in this domain. Our aim is to ensure that the future of fMRI in communication science is healthy, robust, and rich in variety, by encouraging all researchers in the field to think critically about the methods they use, whether that means adopting new analysis methods that can answer previously unanswerable questions, or adjusting their use of methods they already use to align with the latest recommended practices.

Introduction

Despite having existed for just over a decade, the still-young field of communication neuroscience is establishing itself as a full-fledged domain that complements more traditional communication research. A number of recent developments, including special issues on the topic (Afifi & Floyd, 2014; Weber, 2015a, 2015b) and the creation of an interest group at the International Communication Association that explicitly encompasses neuroscientific methods (namely, the Communication Science and Biology group; see also Floyd & Weber, 2018), are a testament to this emergence. With the field on an upward trajectory, but still small enough to be steerable, this article moves beyond the sorts of gentle introductions and defenses of the approach that have come previously (e.g., Weber, Mangus, & Huskey, 2015; Weber, Fisher, Hopp, & Lonergan, 2017), and instead focuses squarely on the future of the field. We do this in two ways: first, by providing details on some of the functional magnetic resonance imaging (fMRI) analytic techniques or approaches that we anticipate gaining traction in the field within the next several years; and second, by taking stock of new challenges that have been identified in the use of analytical methods in fMRI that are already widespread, with the goal of ensuring that the field is built on a strong foundation that will last well into the future.

We presume that readers are familiar with the basics of fMRI, either because they or their collaborators use it themselves, or from having read the aforementioned introductions. Of course, readers need not be familiar with every method we discuss (in fact, we assume some methods will be

new to many readers). However, we err on the side of providing greater technical detail where possible, so a reader with no background in fMRI (for instance, someone reviewing a paper where one of these methods was employed and looking here for a primer) may be overwhelmed by some of the jargon. Such a reader is strongly encouraged to read the two publications mentioned above before reading the present article (and for those looking for an even more detailed introduction to fMRI, we recommend the text by Ashby, 2011). On the other hand, we hope that even advanced readers will find useful information here—if nothing else, from the parallels we draw with other fields using fMRI. As these readers will surely be aware, cognitive neuroscience is in the midst of a self-correction period, with papers challenging various statistical (e.g., Eklund, Nichols, & Knutsson, 2016; Westfall, Nichols, & Yarkoni, 2017) or methodological (e.g., Carp, 2012; Poldrack et al., 2008) aspects of fMRI analysis, as well as pointing out the widespread lack of power (and therefore, replicability) that plagues most fMRI studies (Button et al., 2013; Cremers, Wager, & Yarkoni, 2017; Szucs & Ioannidis, 2017; Turner, Paul, Miller, & Barbey, 2018), appearing almost monthly. Due to its relatively more recent emergence, as well as the richer (i.e., more naturalistic) stimuli often used, communication neuroscience has the opportunity to avoid many of these problems, and to complement the simpler paradigms common in cognitive neuroscience (Huk, Bonnen, & He, 2018).

The present article is organized into two halves around the two future-focused aspects mentioned above, namely introducing analytical methods which are less widely used presently, or critically examining those that are more established. Each half is further subdivided into three sections, corresponding to three classes of fMRI analysis methods, depending on whether they are designed to assess functional activity, functional connectivity, or inter-individual/-group patterns. Note that although the lesser-known analytical methods are generally newer, we do not mean to imply that they will replace the more established methods. Even within a category, different analytical methods are best for answering different sorts of questions, and we anticipate that all of these analytical methods (as well as many we do not discuss here) will coexist going forward. Instead, our focus is slightly different for each half: for the lesser-known analytical methods, we introduce each method and some of its most noteworthy variants or extensions, and also detail the assumptions and limitations of each method, while for the established analytical methods, we largely forego the introduction and focus on outstanding challenges or questions around the use of each method. We conclude with a summary of where the field currently stands, and where we hope to see it go in the next several years.

Overview of classes of analysis methods

The three classes of analysis methods that are considered here will no doubt be familiar to readers with fMRI experience, but in order to clarify our terminology, we will briefly describe each class here (although we leave a description of the specific methods we consider until later), and will also provide a unifying example from the communication literature to further elucidate what sorts of questions each class might be used in addressing. The first class of analysis methods includes those concerned with *Functional Brain Activity*. This class of analysis methods, which aim to identify *where* (and possibly *when*) activity in the brain occurs, is used in a majority of fMRI papers. It is exemplified by the convolution-based general linear model (GLM; Friston et al., 1994), which, owing to its ubiquity, we do not consider further here (although note that “GLM” is often used in other fields to mean “generalized linear model,” and moreover that many other fMRI analysis methods in fact use general or generalized linear models, but that “GLM” almost always refers to this specific application within the field of fMRI). Frequently, this class of analysis methods is based on *subtraction logic*, i.e., the assumption that cognitive processes summate linearly, and that constructing contrasts between conditions that are thought to differ in terms of a single process of interest will yield an understanding of that process.

The second class of analysis methods includes those based on assessing or interpreting *Functional Connectivity*, rather than mere activity—that is, how different areas in the brain communicate with

one another. The recent rapid growth of this class of analysis methods has largely been a product of a shift in the field of cognitive neuroscience from focusing on specific brain regions, to instead conceptualizing the brain as a set of interacting networks (Friston, 1994; Medaglia, Lynall, & Bassett, 2015; Sporns, 2011). Analytical methods for assessing functional connectivity allow researchers to begin to understand how these brain networks operate. There are a number of analytical methods for assessing functional connectivity—in fact, more than 40 different approaches exist (Wang et al., 2014)—but most of these are ill-suited to the sorts of paradigms used most widely in communication neuroscience.

We label the final set of analysis methods we consider here as *Synthetic Analysis Methods*. These are not tied to a particular procedure or modality, but instead, are higher-order analyses that can be applied across multiple domains. That is, in general, these approaches take the outputs of other analysis methods, including from the classes described above, as their inputs. Because of their focus on extending beyond the organization or function of the brain (which, to be clear, is a focus that can be useful to communication scholars, especially given that studies related to brain function lay the foundation for synthetic analyses) and their potential applicability, the analysis methods in this class may ultimately be of the greatest interest to communication researchers.

Having broadly laid out the classes of analysis methods we consider in this article, we now introduce an example framework that we will use throughout in order to demonstrate how each analysis method can be employed to address communication-relevant questions. Our general theoretical domain will be persuasion, in particular in the context of health communication through use of public service announcement (PSA) campaigns. Because the examples for some analysis methods come from our own studies, we also describe some of the theoretical background that underpins this research program—but of course, this is not an empirical paper, so we will gloss over most of the details.

Example: Using fMRI to study persuasion

The framework in which we consider persuasion comes from several well-known communication theories, chiefly including the limited capacity model of motivated mediated message processing (LC4MP, Lang, 2000; see also Fisher, Huskey, Keene, & Weber, 2018; Fisher, Keene, Huskey, & Weber, 2018), as well as the elaboration likelihood model (ELM, Petty & Cacioppo, 1986), and the activation model of information exposure (AMIE, Donohew, Palmgreen, & Duncan, 1980). These models all make predictions about how various message features should affect individuals' attitudes, beliefs, and behavior (Weber, Westcott-Baker, & Anderson, 2013). Two of these features are the *argument strength* (AS) of a message, and the arousingness of a message, which is often operationalized using *message sensation value* (MSV). One simple way in which messages' putative impact can be operationalized is to ask viewers to provide ratings of *perceived message effectiveness* (PME), although the ideal of course is to observe the actual changes in behavior as a consequence of effective persuasive messages. Although these theories do not make reference to brain regions, several brain regions, including the media prefrontal cortex (mPFC) and the dorsolateral prefrontal cortex (dlPFC) within the middle frontal gyrus (MFG) among others, have been empirically implicated repeatedly in linking persuasive message processing to attitudinal and behavioral change (e.g., Chua, Liberzon, Welsh, & Strecher, 2009; Falk, Cascio, & Coronel, 2015; Huskey, Mangus, Turner, & Weber, 2017; Weber, Huskey, Mangus, Westcott-Baker, & Turner, 2015). By dint of their frequent association with persuasion, these regions will be mentioned several times in our examples throughout the article.

To demonstrate how we will use this study domain to illustrate the analytic methods described below, we provide here an example analysis for each of the three classes described at a high level of abstraction—that is, without specifying any particular analysis at this point. A typical analysis in the *Functional Brain Activity* class might ask the degree to which messages—for instance, in the form of PSAs—with differing AS produce differential activity within the persuasion-related brain regions

mentioned above. An analysis focused on *Functional Connectivity* might hypothesize that ads with high MSV—i.e., those with lots of movement, rapid scene cuts, and so forth—will disrupt processing by entraining activity between areas of the brain involved in visual processing and those involved in allocating attention (i.e., such that the visually evoked activity cascades along to attentional areas and drives activity in those areas in a way that inhibits deeper processing). Finally, a *Synthetic Analysis Method* might entail integrating message features and brain activity in (or connectivity between) key regions implicated in valuation and self-referencing such as the mPFC in order to predict which ads will be most effective in producing behavioral change, and for which individuals. With this exemplar framework set up, we now turn to the analysis methods themselves.

Emerging Analytical Methods

In this section, we introduce three methods—one for each of the categories described above—that, to this point, have been used at most rarely within communication neuroscience. Our aim is to provide readers with an understanding of what the analysis method is meant to accomplish, how it works, where they can find the detailed information needed to actually implement it, and some of the most common ways in which it can go wrong.

Functional brain activity: multi-voxel pattern analysis (MVPA)

Overview

In contrast to most other methods for identifying brain activity such as GLM and intersubject correlation analysis (ISC; described in detail later), which consider each voxel or brain region independently, *multivariate methods* such as multi-voxel pattern analysis (MVPA) are based on combining information across multiple voxels (Davis & Poldrack, 2013). Moreover, the philosophy underlying MVPA is fairly distinct from that underlying the GLM. This is because MVPA grew out of a machine learning framework, and is not based on subtraction logic. Instead, MVPA assumes only that different processes or classes of events should be separable when viewed in an appropriate representational space. In other words, the basic question asked by MVPA is: in a given set of voxels, do the *patterns* of activity for one class of events differ reliably from some other class or classes? Moreover, far from being a replacement for univariate methods, there is evidence that the information revealed by univariate and multivariate analyses may in some cases be orthogonal, which is consistent with the evidence supporting both localized and distributed organizations of function within the brain (Haxby, Connolly, & Guntupalli, 2014; Kanwisher, 2010).

In the context of our running persuasion example, MVPA might be used to understand the role that emotional processing plays in (persuasive) ad perception. For instance, a researcher might take a low-cost (cognitively speaking) real-time measure of ad perception by asking participants to press one button during the ad any time they heard a statement with which they agreed, and a second button whenever they heard a statement with which they disagreed. Using these button presses (possibly combined with content analysis of the ads to align the responses with spoken content), MVPA techniques could be used to discover whether activity in a part of the brain previously demonstrated to be involved in emotional processing differed between agreed- and disagreed-with statements—and, moreover, could even be used to generate estimates of the strength of agreement for each specific statement (by examining the raw classifier outputs).

Implementationally, MVPA starts by generating estimates of the unique brain activity associated with every instance of every class of events of interest for every voxel in the brain. The predominant approaches for generating these estimates are special cases of the GLM, and typically return an estimate of the activity associated with each unique stimulus, as either a scalar value (Mumford, Turner, Ashby, & Poldrack, 2012) or a temporal vector (Turner, Mumford, Poldrack, & Ashby, 2012). That is, returning to our persuasion example, the traditional GLM approach would yield a single regression estimate reflecting the mean activity across all stimuli associated with either

agreement or disagreement. By contrast, instance-wise methods such as MVPA will yield an estimate for every unique presentation—so if there are ten agreed-with statements presented during a scan, there will be ten estimates (per voxel), one for each of those ten statements.

The next step in an MVPA analysis is to use these instance-wise estimates to train a pattern classifier (e.g., naïve Bayes, Fisher's linear discriminant, or support-vector machine/SVM) to distinguish between two or more classes (see, e.g., Misaki, Kim, Bandettini, & Kriegeskorte, 2010; Mitchell et al., 2004; Pereira, Mitchell, & Botvinick, 2009). Briefly, voxels form the dimensions of a representational space, and each stimulus instance is a point in this (typically high-dimensional) space. The pattern classifier attempts to find a boundary (called a *hyperplane* for dimensionality greater than three) that separates the stimulus instances of one class from those of the other(s). The measure of interest in these analyses is generally whether out-of-sample classification accuracy is above chance (which can be assessed using the PyMVPA toolbox, which can also be used to fit these models, among other utilities; Hanke et al., 2009). For instance, in our example of discriminating agreed- and disagreed-with statements in an emotional processing region, above-chance accuracy would indicate that the representation of these classes of statements differed within this region. Although this is a similar interpretation as might be obtained from a contrast using the GLM, MVPA can detect discriminative information in situations that thwart the GLM, for instance when information is encoded in a pattern across multiple voxels within the same region and is therefore lost in a subtraction-logic analysis, or when various assumptions of the GLM are violated, (e.g., assumptions of homogeneous variance or linearity of response, depending on choice of classifier).

Assumptions and pitfalls

Probably the most substantial pitfall of these methods is their power and flexibility. Without appropriate cross-validation—that is, testing the classifier on observations that have not been used for the model's estimation—classification accuracy is typically near 100%. Given the number of hyperparameters involved in many classifiers (e.g., bias and slack parameters using a SVM classifier), researchers must be cautious to avoid circularity and overfitting. In fact, overfitting is such a serious issue in pattern classification research that much of the work in the engineering and statistical fields responsible for developing these methods goes towards techniques for minimizing overfitting. In the extreme, a whole-brain classifier will have on the order of 100,000 dimensions, which are used to predict on the order of 100 items. As described above, while the technical details differ, this is analogous to fitting a regression on 100 observations using 100,000 predictor variables; except in very special cases, a perfect fit is mathematically guaranteed. Researchers should therefore be cautious in eagerly adopting some of the most powerful classification techniques—nonlinear classifiers, boosting/bagging, and the like—which are developed for the purposes of maximizing prediction accuracy while sacrificing interpretability (and which are often designed for classification problems with substantially different features than MVPA).

Another problem involves feature selection (i.e., choosing which voxels are included in the MVPA). Rather than select features that are discriminative in a univariate analysis or that have strong weights in an initial MPVA, features should be chosen in a way that is blind to the outcome, which typically means either using *a priori* ROIs or carrying out a whole-brain analysis (potentially with dimensionality reduction such as PCA to partially ameliorate the “curse of dimensionality”). Taken together, the current best-practice pipeline involves using a classifier like a linear SVM, running hold-one-run-out cross-validation (which necessitates the collection of multiple runs; see also Mumford, Davis, & Poldrack, 2014), fixing hyperparameters at reasonable values (referring to the literature for guidance), choosing features in a non-circular way, and assessing classification accuracy significance using nonparametric tests (e.g., a permutation test, Golland & Fischl, 2003).

Variations and extensions

One popular method closely related to MVPA is representational similarity analysis (RSA; Kriegeskorte, Mur, & Bandettini, 2008). Like MVPA, RSA starts with per-instance activity estimates.

However, rather than fitting a classifier to find a separating hyperplane as in MVPA, RSA embeds all of these estimates in a similarity (or equivalently, distance or dissimilarity) matrix. That is, the activation pattern within an ROI for each stimulus instance is statistically compared with the pattern for every other instance (exactly as in the first step of multidimensional scaling). The real power of RSA comes from the next stage of analysis, in which this neural similarity matrix is compared against similarity matrices derived from other sources, either from other brain activity data (for instance, from implanted electrodes), or more relevant for communication scholars, from a theory or computational model. For instance, returning to the running example of persuasion, a published study used RSA to relate several different features of anti-smoking ads to behavioral outcomes (Pegors, Tompson, O'Donnell, & Falk, 2017). In particular, Pegors and colleagues compared the matrix of how similar the brain activity pattern evoked in mPFC by a series of graphic ads was to predictions made on the basis of features such as valence (which would predict high similarity for pairs of ads with a similar valence and low similarity for pairs with opposite valences) and social content (ads with social content would be similar to one another, as would those without, with low similarity between the two categories). They found a strong relationship between how well the social feature predicted the observed neural similarity within each individual, and the individual's behavioral change. In contrast, returning to the idea that multivariate and univariate methods provide complementary information, univariate measures in this same brain region linked the amount of negative valence in the ads to behavioral change.

Another related approach involves the use of *encoding models* (Kay, Naselaris, Prenger, & Gallant, 2008; Naselaris, Kay, Nishimoto, & Gallant, 2011). Briefly, encoding models begin with the collection of large amounts of data, typically during quasi-naturalistic viewing conditions. The stimuli are themselves characterized along a large number of dimensions—everything from low-level decomposition of the images themselves using Gabor filters, to semantic- or content-level representations (Huth, Nishimoto, Vu, & Gallant, 2012)—and these characterizations are used to identify which voxels respond to which dimensions. Encoding models have been employed extensively with stimuli that are very relevant to communication researchers, such as movie clips (Nishimoto et al., 2011). Indeed, by reversing encoding models, it is (theoretically) possible to *decode* brain activity, a process often referred to by the popular press as “mind-reading.” Such claims are now somewhat premature, but this is an area of active development.

Functional brain connectivity: graph theoretic analyses

Overview

Graph theory is not itself a method of measuring functional connectivity, but instead, represents an approach for interpreting the results of functional connectivity analyses. Graph theory is a field of pure mathematics that is used to describe the properties of various classes of networks. Indeed, graph theory is relatively widely used in communication, particularly in the study of social networks (Bond et al., 2012; Huskey, Wilcox, & Weber, 2018; Eagle, Pentland, & Lazer, 2009; Monge & Contractor, 2003; Weber, Alicea, Huskey, & Mathiak, 2018). In the case of functional connectivity, graph theoretic methods can be applied to abstract results away from, say, the connection between a single pair of brain regions, to the system-level properties of an entire network. For example, a basic functional connectivity analysis might report how the connectivity between the mPFC and ventral striatum (which is involved in reward processing) changes as a function of subjective-valuation during media exposure. A graph theoretic approach would examine the entire network to which those brain regions belong (for instance, an affective or valuation network) and describe how efficiently information can flow through the entire network as a function of attention.

For interested readers, we recommend Newman (2010) as an introductory text (for more technical treatments, see Fornito, Zalesky, & Sporns, 2016; or Bassett & Sporns, 2017). In general, graph theoretic analyses begin with a connectivity matrix quantifying the strength of connection between each ROI (or voxel) and every other ROI. In the terminology of graph theory, this is known as an *adjacency matrix*.

Each region that comprises this matrix (i.e., the entities that define the rows/columns of the matrix) are referred to as *nodes*, while the connectivity values themselves reflect the strengths of the *edges* between these nodes. Often, the adjacency matrix (or *graph*) is thresholded such that connections below a certain value are zeroed out, while those above are left intact (a *weighted graph*) or binarized (an *unweighted graph*). On the basis of the resulting *sparse graph*, a multitude of summary statistics can be computed to capture various aspects of the organization of the graph. As a concrete example, one widely used graph theoretic measure is the characteristic path length; this measure begins by computing the minimum number of nodes that must be traversed to get from each node to each other node, and then averaging across all pairs of nodes. This measure is abstracted away from the spatial location of nodes, but indexes something about how efficiently information might travel within this network.

A published example in the domain of persuasion comes from Cooper, Bassett, & Falk (2017). In this work, the authors examined a property related to the cohesiveness of the regions belonging to a set of putative *a priori* networks. After carrying out their graph theoretic analysis from data collected during an ad-rating task using graphic anti-smoking ads, they had a scalar value for each of four networks of interest (such as the *frontoparietal control network*, which is implicated in cognitive control and higher-order goal directed behavior) for every individual, which they could then relate to various measures of behavior change. Their results showed that this graph theoretic property in several networks, including the frontoparietal control network (which includes the dLPFC) was related to individuals' intention to quit, and also demonstrated that a related property in mPFC predicted actual behavior change.

Assumptions and pitfalls

As with other methods discussed herein, the time series data used in a graph-theoretic analysis must undergo fairly standard pre-processing. However, a number of additional decisions need to be made. First among them is how to define the nodes of the network. Several anatomical atlases exist (e.g., Harvard-Oxford), but these are highly variable in node-size and the nodes themselves are often involved across a number of psychological processes and brain networks (that is, the nodes are not coherent with respect to their cognitive or neural functions; but see Van Essen, 2005). As a solution, most researchers in this area use functionally defined atlases, which are well-suited for group analyses and are functionally coherent. However, the decision on which functionally defined atlas to use is non-obvious, as options exist that parcellate the brain into, for example, 50, 90, 164, 200, 264, or 1000 nodes (Craddock, James, Holtzheimer, Hu, & Mayberg, 2012; Power et al., 2011; Shirer, Ryali, Rykhlevskaia, Menon, & Greicius, 2012). In addition to these functionally defined networks, communication scientists may have specific hypotheses about the networks involved and therefore wish to use a set of network-specific nodes for their analysis, typically as defined using meta-analyses (e.g., Yarkoni, Poldrack, Nichols, Van Essen, & Wager, 2011).

A second issue comes with defining what constitutes a meaningful edge. As discussed above, most analyses require a sparse network graph, which is usually accomplished by setting a threshold value for connections between nodes in the network. While no standard has been established, a common convention for whole-brain parcellations is to retain a proportion of the connections as determined by connectivity strength. Additionally, while weighted directional graphs are possible in principle (i.e., edges have a directionality and magnitude specifying how information flows *from* one node *to* another, rather than just *between* the two), most current graph-theoretic applications to fMRI data are usually accomplished on binary undirected graphs. In such instances, important information about magnitude and directionality of connectivity between nodes in a network is lost.

A final consideration on statistical inference is worth noting. Many graph-theoretic measures result in scalar values (e.g., a graph's characteristic path length is a single number). The most common approach (although far from the only one; see Butts, 2009) to combining across individuals is to calculate graph-theoretic statistics on a group-level adjacency matrix (since there is likely to be considerable noise in each individual subject's sparse graph). Under such circumstances, the question of what constitutes an appropriate null-model for a group-level graph metric is non-trivial. The

current best-practice is to generate random models where edges are randomly determined while preserving the overall size and degree of the network (Maslov & Sneppen, 2002). However, there are a variety of solutions that are more or less appropriate for different analyses, and no current gold standard exists (see, e.g., Rubinov & Sporns, 2010; Sizemore & Bassett, 2017; Zalesky, Fornito, & Bullmore, 2012). The current best-practice pipeline is to pre-process fMRI data without spatial smoothing (Alakörkkö, Saarimäki, Glerean, Saramäki, & Korhonen, 2017), form a group-level adjacency matrix from the Pearson correlation coefficients of the detrended neural timeseries from each ROI in a functionally defined atlas (e.g., Power et al., 2011), apply multiple thresholds to this matrix (as graph theoretic statistics can vary by threshold), calculate the relevant graph theoretic statistic for each threshold, calculate the same statistic on null adjacency matrices at each threshold (e.g., Hirschberger, Qi, & Steuer, 2007; Maslov & Sneppen, 2002), and evaluate significance using non-parametric methods (Snijders & Borgatti, 1999) while being sure to account for inflated family-wise error rates (e.g., by applying a false discovery rate correction).¹

Variations and extensions

The procedures described above characterize static graphs, while recent developments have extended these approaches to investigate network dynamics (Sizemore & Bassett, 2017; Telesford et al., 2016)—in fact, the graph-theoretic example from Cooper et al. (2017) given above used just such a dynamic approach. The main advance in using dynamic analyses is the ability to characterize changes in network properties over time. This has clear implications for communication scientists, who are often interested in the ways in which communication dynamics shape behavior and message processing. More broadly, these dynamic approaches may drive considerable advances across a number of core communication questions (Falk & Bassett, 2017), although possibly because they multiply the number of tests or subdivide the data, there are also questions regarding their reliability relative to static methods (Zhang, Baum, Adduru, Biswal, & Michael, 2018).

Synthetic analysis methods: individual differences

Overview

Generally, the above-discussed methods treat individuals as interchangeable (although not always; for instance, both published persuasion examples described above made predictions on a per-individual basis). In some cases, this assumption is relatively defensible: for instance, aside from differences in anatomical structure, brain activity associated with making a simple motor action is fairly consistent across individuals. However, particularly in the case of naturalistic, higher-order stimuli of the sort that are often employed in communication neuroscience, any two individuals' brain responses are likely to differ substantially. Indeed, even in the relatively controlled domain of simplistic cognitive tasks, previous research has shown that differences between individuals performing the same task surpass the differences observed within an individual performing different tasks (Miller et al., 2009); in a more communication-relevant paradigm in which participants listened to long-form stories, activity patterns unique to the individual explained roughly twice as much variance as activity patterns common across individuals (both of which swamped the variance that could be attributed to a specific stimulus; Huth, de Heer, Griffiths, Theunissen, & Gallant, 2016). However, these differences are also not random noise, but are instead explainable in ways that illuminate the underlying stimulus feature space (Van Horn, Grafton, & Miller, 2008) and can speak to theoretical questions (Vogel & Awh, 2008).

The simplest and most widely used approach to studying individual differences in fMRI typically entails relating a single neural measure (e.g., average brain activity in a given brain region) to a single individual difference factor (e.g., educational attainment) by means of a simple correlation (Dubois

¹For example code using this pipeline, see the Open Science Framework page associated with Huskey, Wilcox, & Weber (in press) <https://osf.io/uf5zw/>.

& Adolphs, 2016), and can potentially be extended to an applied or translational context (Gabrieli, Ghosh, & Whitfield-Gabrieli, 2015; Seghier & Price, 2018). We have already described two examples of individual-focused analyses in the context of persuasion; Pegors et al. (2017) used standard linear regression models to relate their RSA-derived measures to behavioral outcomes, while Cooper et al. (2018) related their graph-theoretic measures to behavioral intentions and outcomes using robust regression.

Assumptions and pitfalls

Because individual differences approaches are fundamentally correlational, the standard qualifications regarding correlational research apply: in particular, ascribing a relationship to any measured variable (implicitly, in favor of any other unmeasured variable) is inappropriate, and causal claims need to be made cautiously, if at all. Additionally, the simple approach described above suffers from two further limitations. First, it is not designed to accommodate inherently multivariate data, for instance, brain activity maps or activity timeseries, for example in response to an anti-drug PSA. Second, it is often prone to abuse by way of (uncorrected) multiple comparisons, for instance when running a separate correlation or regression for each outcome measure or using the neural measures from each of a series of ROIs. Ultimately, this approach is limited when it attempts to reduce an individual to one or a small number of features, and ignores all of the other features that might be covariates or confounds. Best practices for this approach include ensuring a sufficient sample size, controlling for multiple comparisons appropriately (preferably by including all measures of interest in a single model—if degrees of freedom are an issue, running multiple models is not an appropriate solution), avoiding circularity in reporting statistics related to the magnitude or significance of the effect (e.g., Vul, Harris, Winkielman, & Pashler, 2009; although cf. important counterpoints to that report, e.g., Yarkoni, 2009, or Lieberman, Berkman, & Wagner, 2009), and limiting conclusions regarding the sensitivity and specificity of any putative relationships to the scope supported by the specific analysis conducted.

Variations and extensions

There is a second approach to studying individual differences in fMRI, which is essentially isomorphic with the RSA approach, except that whereas RSA operates on matrices of similarities between stimuli, the *similarity-based* approach to individual differences operates on matrices of similarities between participants (Miller, Donovan, Bennett, Aminoff, & Mayer, 2012; Miller et al., 2009). For instance, given spatial patterns of brain activity associated with a particular process (say, persuasive message processing) within a particular brain region for each of a number of individual participants, the similarity-based approach would first compute the similarities among these patterns of brain activity to construct an individual-to-individual similarity matrix. Then, these individuals would be characterized in terms of the similarities along a number of other dimensions, such as trait empathy, need for cognition, personality traits, and so forth, in each case yielding a new similarity matrix specific to that domain or measure. Finally, the neural similarity matrix would be regressed on the full complement of similarity matrices for the other dimensions, and the results would indicate how all of these individual difference factors relate to brain activity similarity (see also multivariate distance matrix regression, Anderson, 2001). Typically, to assess significance, variables (or groups of variables) are added to the model sequentially or in a round-robin procedure in order to be able to ascribe unique explanatory power to each variable relative to a set of logically “lower-order” or nuisance variables.

As with the distinction between RSA and MVPA, this approach to individual difference analyses can be used to address the same sorts of questions as the simpler approach described above. However, because this approach abstracts representations away from the original units (whether that is neural activity, graph theoretic measures, personality scale scores, or behavioral performance) into a similarity space, it simplifies testing for isomorphisms that may be impossible to detect otherwise. For example, considering the many ways in which mPFC has been implicated in behavior

change, an individual difference analysis of this sort might start with similarities between individuals in terms of various task-independent (i.e., trait) measures, such as need for cognition or personality, and test how similarity on various outcome or performance measures (such as behavior change, intention to quit, or ratings of the ads) was mediated by similarity across a number of neural measures for the mPFC (e.g., how similar RSA matrices were between individuals, how similar the mean pattern of brain activity was for a particular contrast, or how similar the brainwide connectivity pattern with mPFC was). Combining across all of these domains and measures, many of which are highly multivariate (such as the brainwide connectivity pattern), would be impossible using the simpler approach, but is natural using this approach.

Challenges with established fMRI analyses in communication science

As we have previously mentioned, there are a number of fMRI analyses that are already fairly commonly used in communication neuroscience, and with which we presume most researchers working in this area are familiar. However, even though these analytical procedures are more widespread in the field than the new wave of analyses we described above, this does not mean that there is a uniform process for applying each, nor that there are no outstanding statistical questions in their application. We emphasize that this is not a criticism of any of these analyses, but a call for researchers to continue grappling with these difficult questions, for the good of the field. We argue that such critical thinking is essential for the survival of these analytical approaches in communication neuroscience—and, indeed, perhaps for the long-term survival of communication neuroscience itself.

Functional brain activity: intersubject correlation analysis

Intersubject Correlation (ISC; Hasson, Nir, Levy, Fuhrmann, & Malach, 2004) is an analytical method for detecting brain activity that, unlike the GLM or MVPA/RSA, requires virtually no knowledge of stimulus timing (see also Hejnar, Kiehl, & Calhoun, 2007). The basic principle underlying this method is simple: the degree to which the temporal pattern of brain activity in response to a common input is similar across individuals can serve as an index of task-related brain activity. In its simplest form, similarity between individuals' brain responses can be computed by taking the Pearson correlation between the voxel time series data for every pair of participants, and by then aggregating these voxelwise correlations across all pairs of participants. It is this computation from which the method derives its name.

At a minimum, using ISC requires participants to have received the same stimulus. This does not necessarily mean that the entire scan must be identical for every participant—for instance, if participants view a series of 30-sec ads, those ads themselves can appear in a random order for each participant—but it rules out some experimental paradigms, e.g., video-game playing (where each participant will experience a unique stimulus based on his or her actions during the game). More subtly, ISC does not detect neural activity like a GLM analysis. Instead, ISC is focused on characterizing the processes that are common across participants. Thus, if two participants process the same stimulus differently—e.g., given an ad warning about alcohol abuse, one participant might recall fond memories of past parties, while a second participant might be reminded of the loss of a family member to an accident involving a drunk driver—ISC will be unable to detect anything related to those divergent processes (although it will likely still detect common activation in lower-level regions like primary visual cortex).

There are also a number of statistical difficulties associated with ISC. Researchers should be cautious when making parametric assumptions regarding the standard error of the mean ISC (i.e., the mean of all pairwise between-participant similarities) because of severe violations to assumptions of independence: in an ISC matrix, each participant's data has an effect on the pairwise ISCs with all other participants, and this leads to dependencies in this ISC matrix. Even one

outlier can severely bias all other ISCs. Recently, Chen and colleagues published some promising work aiming to solve the problem of how to specify the standard error of an ISC matrix with inherent dependencies, but this solution will need to be validated thoroughly (Chen, Taylor, Shin, Reynolds, & Cox, 2017). Until this novel solution is validated, the state-of-the art for partly overcoming (but not solving) this issue is to generate randomly varying temporal offsets for the empirical time series data, calculate ISCs, and to treat the resulting “random ISCs” as a simulated “null” distribution (for a toolbox see Kauppi, Pajula, & Tohka, 2014; and for a modern example, see Herbec, Kauppi, Jola, Tohka, & Pollick, 2015).

Additionally, it is not straightforward to map ISC results to constituent processes. In some cases, this is not a problem. For example, if two sets of ads which differ along a defined theoretical dimension (e.g. MSV) are analyzed independently and shown to result in consistently different ISCs, then it may be possible to make statements relating that theoretical dimension and the observed patterns of differential ISCs. In contrast, if a researcher presents participants with a single movie clip and analyzes it with ISC to produce a single map, it is not trivial to determine what aspects of the movie are associated with the observed ISCs. Therefore, it is important that ISC analyses are either accompanied by theoretical considerations from which stimulus conditions can be derived, or that ISC analyses are augmented by event related content analyses of the stimulus material (Grall & Schmälzle, 2018; Weber, Ritterfeld, & Mathiak, 2006).

Finally, when ISC analyses refer to ROIs, it is important to consider how to aggregate ISCs within each ROI. There are two procedures: (1) one can compute voxelwise pairwise ISCs and then average those ISCs for all voxels within target ROIs; or (2) one can first average all voxelwise time series within ROIs, i.e., obtain one average time series for each ROI, and then compute ISCs for all target ROIs. The reported ISCs of these two procedures usually differ substantially with higher ISCs in the latter due to reduced noise in the averaged ROI time series (as demonstrated using a different averaging procedure by Schmälzle, Imhof, Grall, Flaisch, & Schupp, 2017).

Functional brain connectivity: psychological-physiological interaction analysis

The most widely used method of computing functional connectivity in communication neuroscience is PPI, which relies on the standard convolution based GLM logic (Huskey, 2016). Like a standard GLM analysis, the first step of a PPI analysis is to construct a model of the brain activity expected to be evoked by the task (e.g., a boxcar model). Simplifying somewhat, when the model is constructed as a contrast, the method is called standard PPI (sPPI); if instead each task component is modeled separately, the method is called generalized PPI (gPPI). Next, preprocessed BOLD activity (from a single voxel, or averaged over voxels) is selected from a given seed ROI. The ideal next step is to *deconvolve* the extracted BOLD timeseries, take a product of this deconvolved timeseries with each of the (unconvolved) task regressors, and then convolve everything to derive the final set of regressors. After these steps, the full model comprises one regressor for the seed BOLD timeseries (the *physiological* regressor), M regressors encoding the task design (where M is the number of classes of events; these are the *psychological* regressors), and M interaction terms between the physiological regressor and each psychological regressor. The resulting parameter estimate for each interaction term represents how connectivity changes as a function of the event or process associated with the corresponding regressor.

As this is a special case of the GLM, all of the assumptions and pitfalls that accompany the GLM apply here. However, PPI analyses bring with them a number of additional assumptions. First, ideally, PPI involves deconvolving BOLD activity; however, the mathematical operation of convolution is not invertible given finite amounts of data, so in practice, deconvolution involves making various assumptions which may vary in their accuracy (Bush & Cisler, 2013; O'Reilly, Woolrich, Behrens, Smith, & Johansen-Berg, 2012). Current evidence suggests that deconvolution is particularly important for event related designs, but less-so for block designs (Gitelman, Penny, Ashburner,

& Friston, 2003); further work is required to determine the degree to which the alternative—taking the product of BOLD activity with convolved task regressors—is suboptimal.

McLaren and colleagues (2012) also show that gPPI approaches yield better model fits compared to sPPI. This has several important implications. Specifically, considering results from a sPPI analysis necessarily introduces two types of error that are more carefully accounted for in gPPI. The first is essentially a Type II error where significant connectivity patterns may not be observed under a sPPI (because formulating the regressor as a contrast imposes a strict hypothesis regarding the difference between conditions). The second, and more worrisome, is a Type I error where connectivity patterns observed under sPPI are potentially false positive results that disappear when a better statistical model (gPPI) is constructed.

A final word of caution. PPI analyses suffer from reduced statistical power (Friston et al., 1997). This is essentially due to multicollinearity issues. The simplest PPI model will have three regressors: (1) task or psychological, (2) physiological, and (3) the interaction term. This interaction term is necessarily correlated with the other two regressors in the model, which ultimately reduces the ability to detect a significant effect. As with any normal GLM, this can be mitigated (with varying degrees of success) by centering the regressors in the model.

Synthetic analysis methods: brain-as-predictor

Our last method is the brain-as-predictor approach. Conceptually, the idea is straightforward. A (relatively) small group of participants is scanned while they are presented with various stimuli (e.g., persuasive anti-smoking ads compared to some non-persuasive baseline). Brain activity or signal changes (treatment vs. control) within certain key brain regions are estimated for each stimulus. Then, these brain activity estimates are used to predict outcomes in large independent samples or even populations. It is not necessarily surprising that brain activity can be used in this manner—but it is quite surprising that it has been demonstrated that brain activity can *improve* prediction above and beyond self-report measures or theoretical predictions made on the basis of stimulus features (e.g., Falk, Berkman, & Lieberman, 2012). This suggests that there is latent information in the brain to which participants do not have conscious access, are not willing to express (e.g., social desirability, counterarguing), or at the very least, which we do not currently have a means of eliciting using more traditional (e.g., paper-and-pencil or physiological) approaches.

One of the first things to consider with brain-as-predictor is that there may be an inherent circularity in using brain data to predict psychological processes and their self-reported or behavioral outcomes if those outcomes have been used to define the contrasts for modeling the brain response in the same sample. Accordingly, real and valid brain-as-predictor analyses must be out-of-sample, that is, brain data in the fMRI sample is used to predict a psychological process and its self-reported or behavioral outcomes in another, independent sample (or use an independent functional localizer task; see description below). In addition, if predicting behavioral outcomes is the primary goal of a brain-as-predictor study (as opposed to identifying potential cognitive mechanisms from a basic science perspective), some additional justification is required before using difficult-to-collect brain data to predict outcomes in independent samples that easier-to-collect self-report data can predict just as accurately, if not better. Thus, any study using a brain-as-predictor approach for the purpose of prediction must confirm that using brain data improves prediction accuracy above and beyond self-report or observational data; otherwise, the emphasis on improving our basic scientific understanding of the brain's role in communication phenomena *per se* should be made explicit.

Once these requirements are met, brain-as-predictor analyses face two crucial challenges: (1) which brain data should be used, and (2) how to model brain data together with self-report and behavioral data across two or more independent samples? Typically, researchers use existing literature to identify structures that are commonly implicated in a given communication-relevant process. Using the most straightforward data-driven alternative—selecting an ROI due to its ability to make out of sample predictions in one study—threatens the generalizability of that ROI for future

studies. This is because ROIs derived from data-driven procedures may capitalize on chance, and therefore may have limited predictive ability in replication studies. Of course, as we have already described, some of the most well-studied ROIs in the area of predicting persuasion outcomes have been validated and replicated across multiple domains and with multiple neural measures and prediction methodologies (e.g., the mPFC). Nonetheless, the safest approach is to use ROIs defined completely independently (e.g., from a functional localizer designed to target a process of interest, e.g., as in Falk et al., 2016; from a data-driven parcellation, e.g., Schaefer et al., 2017; from a meta-analytic tool such as neurosynth, Yarkoni et al., 2011; or based on anatomical features, e.g., Glasser et al., 2016).

Moreover, in addition to choosing the ROI, the choice of brain data entails choosing a valid control condition. Using our running study example, if PSAs varying in persuasiveness are used to predict behavior change, then the best control condition for the neural predictors is a PSA from which (only) the persuasive content in question has been removed. This is not a trivial procedure. For instance, researchers could use modified PSAs in which the relevant message has been removed or altered, or use the same PSAs presented in a scrambled version (i.e., removing any meaning but keeping sound and brightness amplitudes constant). The use of a task-free “baseline” (i.e., a condition in which participants focus on a black screen with a centered cross) is common, but may fail to properly control for various confounds, and also raises the possibility of baseline differences (e.g., between individuals) adding noise and degrading prediction quality. Of course, one could argue that if a brain-as-predictor analysis has produced out-of-sample predictions of message effectiveness and behavior change with increased accuracy over self-reports, then the baseline or confound issue was obviously not a major problem. We agree with this counter argument if the practical application of brain-as-predictor analyses is at the center of a study. However, for the further development of persuasion theory, that is, for enhancing our understanding of why exactly certain messages work and others do not, a carefully designed control condition that can isolate explanations to particular message features is important. To be clear, such a lack of a proper control or baseline condition is not unique to or inherent in brain-as-predictor analyses, but the role of a proper baseline has perhaps previously been articulated less clearly than for other analysis methods.

Next, researchers using the brain-as-predictor approach need to use brain and self-report data in the fMRI sample to predict self-report or behavioral data in another sample. The current practice to combine data from heterogeneous datasets is to rank order all available data and to use an ordinal regression model for prediction. Despite this general practice in virtually all published brain-as-predictor studies, it is not obvious that this procedure is appropriate under all circumstances. Dependent on the data’s original distributions, rank transformations can artificially inflate or deflate prediction accuracies (because rank transformations are inherently nonlinear, which arbitrarily changes the data’s covariance structure; Akritas, 1991; Saltelli & Sobol, 1995; Seaman, Algina, & Olejnik, 1985). Thus, albeit rarely done, it is good practice to report the data’s distribution characteristics (especially variance and skew) and to compare the ranked and unranked prediction models before making statements about the true predictive power of brain-as-predictor analyses.

Finally, brain-as-predictor still requires classic forward inference and careful experimental manipulation in order to determine what exact role any given brain area plays in a particular phenomenon. That is, the observation that activity in a brain area predicts outcomes in an independent sample above chance levels certainly demonstrates the sufficiency of that activity in making such predictions, but does not shed light on the necessity of that brain area (i.e., it could simply be upstream or downstream of some region that is actually causally involved); likewise, beyond vague process-level descriptions like “persuasive message processing,” determining the mechanism requires going beyond simple message-level prediction (e.g., as exemplified by the sort of more model-driven analyses featured in more recent work such as Pegors et al., 2017). This limitation is in some ways related to the reverse inference fallacy (the mistaken belief that observing activity in brain area X during putative cognitive process Y means that X *does* Y) that is common across all of the methods described here (as well as most other fMRI analysis methods; Poldrack, 2006). However, it is

exacerbated when the process in question is not in fact a process, but a phenomenon that results from potentially dozens of interacting processes. Notwithstanding, converging evidence can certainly shed light on the role of well-studied areas, as the examples related to the mPFC and persuasion that we cite above show.

Conclusion

The use of fMRI is fairly new in communication, but as it becomes more common, the range of analytic methods employed grows as well. For researchers who may want to begin carrying out this type of research themselves, it can be overwhelming to keep track of the goals and limitations of all of these methods. Our aim here is to give these researchers a jumpstart and some understanding of these goals and limitations, and also to provide canonical references for those interested in learning more. Likewise, even for researchers who work with some of these methods routinely, it can be helpful to be reminded of some of the ways in which analytical methods for fMRI data can fail, particularly when such failures are hard to detect. Therefore, although we avoided highly technical detail, we tried to highlight challenges beyond simple “rookie mistakes.”

By making these analytical methods somewhat more accessible, and offering examples of how each method has been used in communication research previously (or how it might be used in the future in the case of methods that have not yet been adopted), we can increase the pace at which these methods are employed in furthering our understanding of communication phenomena. Moreover, by pointing out unanswered challenges in the technical application of some of the more widespread analysis techniques within the field, we specify the work that still needs to be done, and highlight the caution required in interpreting otherwise dazzling results. While it cannot (and should not be expected to) supplant other more traditional approaches to communication research, we believe strongly in the promise of fMRI to advance the field, and to help catalyze a new era of research.

References

- Afifi, T., & Floyd, K. (2014). Biology and physiology in communication [special issue]. *Communication Monographs*, 82(1), 1–178.
- Akritis, M. G. (1991). Limitations of the rank transform procedure: A study of repeated measures designs, Part I. *Journal of the American Statistical Association*, 86, 457–460. doi:10.1080/01621459.1991.10475066
- Alakörkkö, T., Saarimäki, H., Glerean, E., Saramäki, J., & Korhonen, O. (2017). Effects of spatial smoothing on functional brain networks. *European Journal of Neuroscience*, 46(9), 2471–2480. doi:10.1111/ejn.2017.46.issue-9
- Anderson, M. J. (2001). A new method for non-parametric multivariate analysis of variance. *Austral Ecology*, 26, 32–46.
- Ashby, F. G. (2011). *Statistical analysis of fMRI data*. Cambridge, MA: MIT press.
- Bassett, D. S., & Sporns, O. (2017). Network neuroscience. *Nature Neuroscience*, 20, 353–364. doi:10.1038/nn.4465
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415), 295–298.
- Bush, K., & Cisler, J. (2013). Decoding neural events from fMRI BOLD signal: A comparison of existing approaches and development of a new algorithm. *Magnetic Resonance Imaging*, 31, 976–989. doi:10.1016/j.mri.2013.03.015
- Butts, C. T. (2009). Revisiting the foundations of network analysis. *Science*, 325, 414–416. doi:10.1126/science.1171022
- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., & Munafò, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365.
- Carp, J. (2012). The secret lives of experiments: Methods reporting in the fMRI literature. *Neuroimage*, 63, 289–300. doi:10.1016/j.neuroimage.2012.06.037
- Chen, G., Taylor, P. A., Shin, Y.-W., Reynolds, R. C., & Cox, R. W. (2017). Untangling the relatedness among correlations, part II: Inter-subject correlation group analysis through linear mixed-effects modeling. *Neuroimage*, 147, 825–840. doi:10.1016/j.neuroimage.2016.08.029
- Chua, H. F., Liberzon, I., Welsh, R. C., & Strecher, V. J. (2009). Neural correlates of message tailoring and self-relatedness in smoking cessation programming. *Biological Psychiatry*, 65(2), 165–168. doi:10.1016/j.biopsych.2008.08.030

- Cooper, N., Bassett, D. S., & Falk, E. B. (2017). Coherent activity between brain regions that code for value is linked to the malleability of human behavior. *Scientific Reports*, 7, 43250. doi:[10.1038/srep43250](https://doi.org/10.1038/srep43250)
- Craddock, R. C., James, G. A., Holtzheimer, P. E., Hu, X. P., & Mayberg, H. S. (2012). A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Human Brain Mapping*, 33, 1914–1928. doi:[10.1002/hbm.v33.8](https://doi.org/10.1002/hbm.v33.8)
- Cremers, H. R., Wager, T. D., & Yarkoni, T. (2017). The relation between statistical power and inference in fMRI. *PLOS One*, 12, e0184923. doi:[10.1371/journal.pone.0184923](https://doi.org/10.1371/journal.pone.0184923)
- Donohew, L., Palmgreen, P., & Duncan, J. (1980). An activation model of information exposure. *Communications Monographs*, 47(4), 295–303.
- Davis, T., & Poldrack, R. A. (2013). Measuring neural representations with fMRI: Practices and pitfalls. *Annals of the New York Academy of Sciences*, 1296, 108–134. doi:[10.1111/nyas.12156](https://doi.org/10.1111/nyas.12156)
- Dubois, J., & Adolphs, R. (2016). Building a science of individual differences from fMRI. *Trends in Cognitive Sciences*, 20, 425–443. doi:[10.1016/j.tics.2016.03.014](https://doi.org/10.1016/j.tics.2016.03.014)
- Eagle, N., Pentland, A. S., & Lazer, D. (2009). Inferring friendship network structure by using mobile phone data. *Proceedings of the national academy of sciences*, 106(36), 15274–15278.
- Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences*, 113, 7900–7905.
- Falk, E. B., & Bassett, D. S. (2017). Brain and social networks: Fundamental building blocks of human experience. *Trends in Cognitive Sciences*, 21, 674–690. doi:[10.1016/j.tics.2017.06.009](https://doi.org/10.1016/j.tics.2017.06.009)
- Falk, E. B., Berkman, E. T., & Lieberman, M. D. (2012). From neural responses to population behavior: Neural focus group predicts population-level media effects. *Psychological Science*, 23, 439–445. doi:[10.1177/0956797611434964](https://doi.org/10.1177/0956797611434964)
- Falk, E. B., Cascio, C. N., & Coronel, J. C. (2015). Neural prediction of communication-relevant outcomes. *Communication Methods and Measures*, 9, 30–54. doi:[10.1080/19312458.2014.999750](https://doi.org/10.1080/19312458.2014.999750)
- Falk, E. B., O'Donnell, M. B., Tompson, S., Gonzalez, R., Dal Cin, S. D., Strecher, V., ... An, L. (2016). Functional brain imaging predicts public health campaign success. *Social Cognitive and Affective Neuroscience*, 11(2), 204–214. doi:[10.1093/scan/nsv108](https://doi.org/10.1093/scan/nsv108)
- Fisher, J., Huskey, R., Keene, J., & Weber, R. (2018). *The limited capacity model of motivated mediated message processing: Looking to the future*. Annals of the International Communication Association. Manuscript submitted for publication.
- Floyd, K., & Weber, R. (Eds.). (2018). *The handbook of communication science and biology*. International communication association handbook series. New York, NY: Routledge.
- Fornito, A., Zalesky, A., & Sporns, O. (2016). *Fundamentals of brain network analysis*. San Diego, CA: Academic Press.
- Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, 2, 56–78. doi:[10.1002/hbm.v2:1/2](https://doi.org/10.1002/hbm.v2:1/2)
- Friston, K. J., Buechel, C., Fink, G. R., Morris, J., Rolls, E., & Dolan, R. J. (1997). Psychophysiological and modulatory interactions in neuroimaging. *Neuroimage*, 6, 218–229. doi:[10.1006/nimg.1997.0291](https://doi.org/10.1006/nimg.1997.0291)
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. P., Frith, C. D., & Frackowiak, R. S. (1994). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2, 189–210. doi:[10.1002/hbm.460020402](https://doi.org/10.1002/hbm.460020402)
- Gabrieli, J. D., Ghosh, S. S., & Whitfield-Gabrieli, S. (2015). Prediction as a humanitarian and pragmatic contribution from human cognitive neuroscience. *Neuron*, 85(1), 11–26.
- Gitelman, D. R., Penny, W. D., Ashburner, J., & Friston, K. J. (2003). Modeling regional and psychophysiological interactions in fMRI: The importance of hemodynamic deconvolution. *Neuroimage*, 19, 200–207. doi:[10.1016/S1053-8119\(03\)00058-2](https://doi.org/10.1016/S1053-8119(03)00058-2)
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., ... Smith, S. M. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615), 171–178. doi:[10.1038/nature18933](https://doi.org/10.1038/nature18933)
- Golland, P., & Fischl, B. (2003). *Permutation tests for classification: Towards statistical significance in image-based studies*. In Proceedings from IPMI: Biennial International Conference on Information Processing in Medical Imaging (pp. 330–341). Berlin: Springer.
- Grall, C., & Schmälzle, R. (2018, May). *From pixels to media effects: A study of 600 brains watching a suspenseful movie*. Paper presented at the 68th Annual Conference of the International Communication Association, Prague, Czech Republic.
- Hanke, M., Halchenko, Y. O., Sederberg, P. B., Hanson, S. J., Haxby, J. V., & Pollmann, S. (2009). PyMVPA: A Python toolbox for multivariate pattern analysis of fMRI data. *Neuroinformatics*, 7, 37–53. doi:[10.1007/s12021-008-9041-y](https://doi.org/10.1007/s12021-008-9041-y)
- Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science*, 303, 1634–1640. doi:[10.1126/science.1089506](https://doi.org/10.1126/science.1089506)
- Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding neural representational spaces using multivariate pattern analysis. *Annual Review of Neuroscience*, 37, 435–456. doi:[10.1146/annurev-neuro-062012-170325](https://doi.org/10.1146/annurev-neuro-062012-170325)
- Hejnar, M. P., Kiehl, K. A., & Calhoun, V. D. (2007). Interparticipant correlations: A model free FMRI analysis technique. *Human Brain Mapping*, 28, 860–867. doi:[10.1002/\(ISSN\)1097-0193](https://doi.org/10.1002/(ISSN)1097-0193)

- Herbec, A., Kauppi, J. P., Jola, C., Tohka, J., & Pollick, F. E. (2015). Differences in fMRI intersubject correlation while viewing unedited and edited videos of dance performance. *Cortex*, 71, 341–348. doi:10.1016/j.cortex.2015.06.026
- Hirschberger, M., Qi, Y., & Steuer, R. E. (2007). Randomly generating portfolio-selection covariance matrices with specified distributional characteristics. *European Journal of Operational Research*, 177(3), 1610–1625. doi:10.1016/j.ejor.2005.10.014
- Huk, A., Bonnen, K., & He, B. J. (2018). Beyond trial-based paradigms: Continuous behavior, ongoing neural activity, and natural stimuli. *Journal of Neuroscience*, 38, 7551–7558. doi:10.1523/JNEUROSCI.0052-18.2018
- Huskey, R., Wilcox, S., & Weber, R. (in press). Network neuroscience reveals distinct neuromarkers of flow during media use. *Journal of Communication*.
- Huskey, R. (2016). Beyond blobology: Using psychophysiological interaction analyses to investigate the neural basis of human communication phenomena. In S. Kubitschko & A. Kaun (Eds.), *Innovative methods in media and communication research* (Vol. 1, pp. 123–138). London, UK: Palgrave Macmillan. Advance online publication
- Huskey, R., Mangus, J. M., Turner, B. O., & Weber, R. (2017). The persuasion network is modulated by drug-use risk and predicts anti-drug message effectiveness. *Social Cognitive and Affective Neuroscience*, 12, 1902–1915. doi:10.1093/scan/nsx126
- Huskey, R., Wilcox, S., & Weber, R. (2018). *Network neuroscience reveals distinct neuromarkers of flow during media use*. Manuscript submitted for publication.
- Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E., & Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532, 453–458. doi:10.1038/nature17637
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron*, 76, 1210–1224. doi:10.1016/j.neuron.2012.10.014
- Kanwisher, N. (2010). Functional specificity in the human brain: A window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences*, 107, 11163–11170.
- Kauppi, J.-P., Pajula, J., & Tohka, J. (2014). A versatile software package for inter-subject correlation based analyses of fMRI. *Frontiers in Neuroinformatics*, 8(2). doi:10.3389/fninf.2014.00002
- Kay, K. N., Naselaris, T., Prenger, R. J., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, 452, 352–355. doi:10.1038/nature06713
- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2, 4.
- Lang, A. (2000). The limited capacity model of mediated message processing. *Journal of communication*, 50(1), 46–70.
- Lieberman, M. D., Berkman, E. T., & Wagner, T. D. (2009). Correlations in social neuroscience aren't voodoo. Commentary on Vul et al. (2009). *Perspectives on Psychological Science*, 4(3), 299–307. doi:10.1111/j.1745-6924.2009.01128.x
- Maslov, S., & Sneppen, K. (2002). Specificity and stability in topology of protein networks. *Science*, 296(5569), 910–913. doi:10.1126/science.1065103
- McLaren, D. G., Ries, M. L., Xu, G., & Johnson, S. C. (2012). A generalized form of context-dependent psychophysiological interactions (gPPI): A comparison to standard approaches. *Neuroimage*, 61, 1277–1286. doi:10.1016/j.neuroimage.2012.03.068
- Medaglia, J. D., Lynall, M. E., & Bassett, D. S. (2015). Cognitive network neuroscience. *Journal of Cognitive Neuroscience*, 27, 1471–1491. doi:10.1162/jocn_a_00810
- Miller, M. B., Donovan, C. L., Bennett, C. M., Aminoff, E. M., & Mayer, R. E. (2012). Individual differences in cognitive style and strategy predict similarities in the patterns of brain activity between individuals. *Neuroimage*, 59, 83–93. doi:10.1016/j.neuroimage.2011.05.060
- Miller, M. B., Donovan, C. L., Van Horn, J. D., German, E., Sokol-Hessner, P., & Wolford, G. L. (2009). Unique and persistent individual patterns of brain activity across different memory retrieval tasks. *Neuroimage*, 48, 625–635. doi:10.1016/j.neuroimage.2009.06.033
- Misaki, M., Kim, Y., Bandettini, P. A., & Kriegeskorte, N. (2010). Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *Neuroimage*, 53, 103–118. doi:10.1016/j.neuroimage.2010.05.051
- Mitchell, T. M., Hutchinson, R., Niculescu, R. S., Pereira, F., Wang, X., Just, M., & Newman, S. (2004). Learning to decode cognitive states from brain images. *Machine Learning*, 57, 145–175. doi:10.1023/B:MACH.0000035475.85309.1b
- Monge, P. R., & Contractor, N. S. (2003). *Theories of communication networks*. Oxford, UK: Oxford University Press.
- Mumford, J. A., Davis, T., & Poldrack, R. A. (2014). The impact of study design on pattern estimation for single-trial multivariate pattern analysis. *Neuroimage*, 103, 130–138. doi:10.1016/j.neuroimage.2014.09.026
- Mumford, J. A., Turner, B. O., Ashby, F. G., & Poldrack, R. A. (2012). Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *Neuroimage*, 59, 2636–2643. doi:10.1016/j.neuroimage.2011.08.076
- Naselaris, T., Kay, K. N., Nishimoto, S., & Gallant, J. L. (2011). Encoding and decoding in fMRI. *Neuroimage*, 56, 400–410. doi:10.1016/j.neuroimage.2010.07.073
- Newman, M. E. J. (2010). *Networks: An Introduction*. Oxford, United Kingdom: Oxford University Press.

- Nishimoto, S., Vu, A. T., Naselaris, T., Benjamini, Y., Yu, B., & Gallant, J. L. (2011). Reconstructing visual experiences from brain activity evoked by natural movies. *Current Biology*, 21, 1641–1646. doi:[10.1016/j.cub.2011.08.031](https://doi.org/10.1016/j.cub.2011.08.031)
- O'Reilly, J. X., Woolrich, M. W., Behrens, T. E. J., Smith, S. M., & Johansen-Berg, H. (2012). Tools of the trade: Psychophysiological interactions and functional connectivity. *Social Cognitive and Affective Neuroscience*, 7, 604–609. doi:[10.1093/scan/nss055](https://doi.org/10.1093/scan/nss055)
- Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In *Communication and persuasion*, (pp. 1–24). New York, NY: Springer.
- Pegors, T. K., Tompson, S., O'Donnell, M. B., & Falk, E. B. (2017). Predicting behavior change from persuasive messages using neural representational similarity and social network analyses. *Neuroimage*, 157, 118–128. doi:[10.1016/j.neuroimage.2017.05.063](https://doi.org/10.1016/j.neuroimage.2017.05.063)
- Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: A tutorial overview. *Neuroimage*, 45, S199–S209. doi:[10.1016/j.neuroimage.2008.11.007](https://doi.org/10.1016/j.neuroimage.2008.11.007)
- Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Sciences*, 10, 59–63. doi:[10.1016/j.tics.2005.12.004](https://doi.org/10.1016/j.tics.2005.12.004)
- Poldrack, R. A., Fletcher, P. C., Henson, R. N., Worsley, K. J., Brett, M., & Nichols, T. E. (2008). Guidelines for reporting an fMRI study. *Neuroimage*, 40, 409–414. doi:[10.1016/j.neuroimage.2007.11.048](https://doi.org/10.1016/j.neuroimage.2007.11.048)
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., ... Petersen, S. E. (2011). Functional network organization of the human brain. *Neuron*, 72, 665–678. doi:[10.1016/j.neuron.2011.09.006](https://doi.org/10.1016/j.neuron.2011.09.006)
- Rubinov, M., & Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage*, 52(3), 1059–1069. doi:[10.1016/j.neuroimage.2009.10.003](https://doi.org/10.1016/j.neuroimage.2009.10.003)
- Saltelli, A., & Sobol, I. M. (1995). About the use of rank transformation in sensitivity analysis of model output. *Reliability Engineering & System Safety*, 50, 225–239. doi:[10.1016/0951-8320\(95\)00099-2](https://doi.org/10.1016/0951-8320(95)00099-2)
- Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X. N., Holmes, A. J., ... Yeo, B. T. (2017). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*. doi:[10.1093/cercor/bhx179](https://doi.org/10.1093/cercor/bhx179)
- Schmälzle, R., Imhof, M. A., Grall, C., Fleisch, T., & Schupp, H. T. (2017). Reliability of fMRI time series: Similarity of neural processing during movie viewing. *bioRxiv*, 158188. doi:[10.1101/158188](https://doi.org/10.1101/158188)
- Seaman, S. L., Algina, J., & Olejnik, S. F. (1985). Type I error probabilities and power of the rank and parametric ANCOVA procedures. *Journal of Educational Statistics*, 10, 345–367. doi:[10.3102/10769986010004345](https://doi.org/10.3102/10769986010004345)
- Seghier, M. L., & Price, C. J. (2018). Interpreting and utilising intersubject variability in brain function. *Trends in Cognitive Sciences*, 22, 517–530. doi:[10.1016/j.tics.2018.03.003](https://doi.org/10.1016/j.tics.2018.03.003)
- Shirer, W. R., Ryali, S., Rykhlevskaia, E., Menon, V., & Greicius, M. D. (2012). Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex*, 22, 158–165. doi:[10.1093/cercor/bhr115](https://doi.org/10.1093/cercor/bhr115)
- Sizemore, A. E., & Bassett, D. S. (2017). Dynamic graph metrics: Tutorial, toolbox, and tale. *NeuroImage*, 180(Part B), 417–427.
- Snijders, T. A. B., & Borgatti, S. P. (1999). Non-parametric standard errors and tests for network statistics. *Connections*, 22(2), 161–170.
- Sporns, O. (2011). *Networks of the brain*. Cambridge, MA: The MIT Press.
- Szucs, D., & Ioannidis, J. P. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biology*, 15, e2000797. doi:[10.1371/journal.pbio.2003145](https://doi.org/10.1371/journal.pbio.2003145)
- Telesford, Q. K., Lynall, M.-E., Vettel, J., Miller, M. B., Grafton, S. T., & Bassett, D. S. (2016). Detection of functional brain network reconfiguration during task-driven cognitive states. *Neuroimage*, 142, 198–210. doi:[10.1016/j.neuroimage.2016.05.078](https://doi.org/10.1016/j.neuroimage.2016.05.078)
- Turner, B. O., Mumford, J. A., Poldrack, R. A., & Ashby, F. G. (2012). Spatiotemporal activity estimation for multivoxel pattern analysis with rapid event-related designs. *Neuroimage*, 62, 1429–1438. doi:[10.1016/j.neuroimage.2012.05.057](https://doi.org/10.1016/j.neuroimage.2012.05.057)
- Turner, B. O., Paul, E. J., Miller, M. B., & Barbey, A. K. (2018). Small sample sizes reduce the replicability of task-based fMRI studies. *Communications Biology*, 1, 62. doi:[10.1038/s42003-018-0073-z](https://doi.org/10.1038/s42003-018-0073-z)
- Van Essen, D. C. (2005). A population-average, landmark- and surface-based (PALS) atlas of human cerebral cortex. *Neuroimage*, 28, 635–662. doi:[10.1016/j.neuroimage.2005.06.058](https://doi.org/10.1016/j.neuroimage.2005.06.058)
- Van Horn, J. D., Grafton, S. T., & Miller, M. B. (2008). Individual variability in brain activity: A nuisance or an opportunity? *Brain Imaging and Behavior*, 2, 327–334. doi:[10.1007/s11682-008-9049-9](https://doi.org/10.1007/s11682-008-9049-9)
- Vogel, E. K., & Awh, E. (2008). How to exploit diversity for scientific gain: Using individual differences to constrain cognitive theory. *Current Directions in Psychological Science*, 17(2), 171–176. doi:[10.1111/j.1467-8721.2008.00569.x](https://doi.org/10.1111/j.1467-8721.2008.00569.x)
- Vul, E., Harris, C., Winkelman, P., & Pashler, H. (2009). Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspectives on Psychological Science*, 4, 274–290. doi:[10.1111/j.1745-6924.2009.01125.x](https://doi.org/10.1111/j.1745-6924.2009.01125.x)
- Wang, H. E., Benar, C. G., Quilichini, P. P., Friston, K. J., Jirsa, V. K., & Bernard, C. (2014). A systematic framework for functional connectivity measures. *Frontiers in Neuroscience*, 8, 1–22. doi:[10.3389/fnins.2014.00405](https://doi.org/10.3389/fnins.2014.00405)
- Weber, R., Westcott-Baker, A., & Anderson, G. (2013). A multilevel analysis of antimarijuana public service announcement effectiveness. *Communication Monographs*, 80(3), 302–330.

- Weber, R. (2015a). Brain, mind, and media: Neuroscience meets media psychology [special issue]. *Journal of Media Psychology*, 25(4), 153–204.
- Weber, R. (2015b). Biology and brains: Methodological innovations in communication science [special issue]. *Communication Methods and Measures*, 9, 1–2. doi:[10.1080/19312458.2014.999755](https://doi.org/10.1080/19312458.2014.999755)
- Weber, R., Alicea, B., Huskey, R., & Mathiak, K. (2018). Network dynamics of attention during a naturalistic behavioral paradigm. *Frontiers in Human Neuroscience*, 12(182), 1–14. doi:[10.3389/fnhum.2018.00182](https://doi.org/10.3389/fnhum.2018.00182)
- Weber, R., Fisher, J. T., Hopp, F. R., & Lonergan, C. (2017). Taking messages into the magnet. Method-theory synergy in communication neuroscience. *Communication Monographs*, 85(1), 81–102. doi:[10.1080/03637751.2017.1395059](https://doi.org/10.1080/03637751.2017.1395059)
- Weber, R., Huskey, R., Mangus, J. M., Westcott-Baker, A., & Turner, B. (2015). Neural predictors of message effectiveness during counterarguing in antidrug campaigns. *Communication Monographs*, 82(1), 4–30. doi:[10.1080/03637751.2014.971414](https://doi.org/10.1080/03637751.2014.971414)
- Weber, R., Mangus, J. M., & Huskey, R. (2015). Brain imaging in communication research: A practical guide to understanding and evaluating fMRI studies. *Communication Methods and Measures*, 9, 5–29. doi:[10.1080/19312458.2014.999754](https://doi.org/10.1080/19312458.2014.999754)
- Weber, R., Ritterfeld, U., & Mathiak, K. (2006). Does playing violent video games induce aggression? Empirical evidence of a functional magnetic resonance imaging study. *Media Psychology*, 8, 39–60. doi:[10.1207/S1532785XMEP0801_4](https://doi.org/10.1207/S1532785XMEP0801_4)
- Westfall, J., Nichols, T. E., & Yarkoni, T. (2017). Fixing the stimulus-as-fixed-effect fallacy in task fMRI [version 2; referees: 1 approved, 2 approved with reservations]. *Wellcome Open Research*, 1–23. doi:[10.12688/wellcomeopenres.10298.2](https://doi.org/10.12688/wellcomeopenres.10298.2)
- Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, 8, 665–670. doi:[10.1038/nmeth.1635](https://doi.org/10.1038/nmeth.1635)
- Zalesky, A., Fornito, A., & Bullmore, E. (2012). On the use of correlation as a measure of network connectivity. *NeuroImage*, 60(4), 2096–2106. doi:[10.1016/j.neuroimage.2012.02.001](https://doi.org/10.1016/j.neuroimage.2012.02.001)
- Zhang, C., Baum, S. A., Adduru, V. R., Biswal, B. B., & Michael, A. M. (2018). Test-retest reliability of dynamic functional connectivity in resting state fMRI. *Neuroimage*. doi:[10.1016/j.neuroimage.2018.08.021](https://doi.org/10.1016/j.neuroimage.2018.08.021)