NeuroImage 146 (2017) 275–287

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Empirical validation of directed functional connectivity

Ravi D. Mill[a], Anto Bagic[b], Andreea Bostan[c,d], Walter Schneider[c,d,e], Michael W. Cole[a,c,d,*]

[a] *Center for Molecular and Behavioral Neuroscience, Rutgers University, USA*
[b] *Department of Neurology, University of Pittsburgh, USA*
[c] *Center for Neuroscience, University of Pittsburgh, USA*
[d] *Center for the Neural Basis of Cognition, University of Pittsburgh, USA*
[e] *Department of Psychology, University of Pittsburgh, USA*

ARTICLE INFO

Keywords:
Functional connectivity
Directed connectivity
Effective connectivity
fMRI
MEG
Memory

ABSTRACT

Mapping directions of influence in the human brain connectome represents the next phase in understanding its functional architecture. However, a host of methodological uncertainties have impeded the application of directed connectivity methods, which have primarily been validated via "ground truth" connectivity patterns embedded in simulated functional MRI (fMRI) and magneto-/electro-encephalography (MEG/EEG) datasets. Such simulations rely on many generative assumptions, and we hence utilized a different strategy involving empirical data in which a ground truth directed connectivity pattern could be anticipated with confidence. Specifically, we exploited the established "sensory reactivation" effect in episodic memory, in which retrieval of sensory information reactivates regions involved in perceiving that sensory modality. Subjects performed a paired associate task in separate fMRI and MEG sessions, in which a ground truth reversal in directed connectivity between auditory and visual sensory regions was instantiated across task conditions. This directed connectivity reversal was successfully recovered across different algorithms, including Granger causality and Bayes network (IMAGES) approaches, and across fMRI ("raw" and deconvolved) and source-modeled MEG. These results extend simulation studies of directed connectivity, and offer practical guidelines for the use of such methods in clarifying causal mechanisms of neural processing.

1. Introduction

The advent of network methods stands as a significant development in human cognitive neuroscience, extending characterization of the function of isolated brain regions to that of connections between regions and large-scale networks of regions (Craddock et al., 2013; Medaglia et al., 2015; Sporns, 2011). To date, this field of network neuroscience has been dominated by methods of "undirected" functional connectivity, which infer whether two brain regions A and B are communicating in some general fashion, as typically revealed by the Pearson's correlation computed between their activity time series (Biswal et al., 1995; Friston et al., 1997). In contrast, "directed" functional connectivity (or "effective" connectivity) methods clarify asymmetries in activity flow that determine whether region A is communicating downstream to region B (connectivity A→B) or vice versa (connectivity B→A). Suggested approaches to analyzing directed connectivity in brain imaging data have included Granger causality (Roebroeck et al., 2005; Seth, 2010), directed coherence (Nolte et al., 2008), dynamic causal modeling (DCM; Friston et al., 2003), linear non-Gaussian (LiNG; Hyvarinen and Smith, 2013), conditional Bayes (Patel et al. 2006) and Bayes network methods (Mumford and Ramsey, 2014).

Whilst the relative capabilities of the above algorithms to map truly "causal" or "effective" connections have been debated (Friston, 2011; Roebroeck et al., 2011), they nonetheless collectively entail a conceptual advance over undirected methods by linking cognitive operations to more precise computational mechanisms. However, widespread application of directed connectivity has been hampered by a number of methodological uncertainties, spanning the choice of imaging modality, directional algorithm, input parameters and preprocessing steps. These uncertainties call for concerted attempts to validate directed connectivity, and it is the aim of the present paper to address this need.

Prior validations have predominantly relied on the recovery of directional patterns embedded in simulated datasets. Much of this work has focused on fMRI given its present popularity, and also as it presents perhaps the clearest challenges to the application of directed connectivity. Specifically, observed BOLD signals are convolved with hemodynamic response functions (HRF), and hence offer an indirect, low-pass filtered, non-linear reflection of neuronal activity. A widely

---

[*] Corresponding author at: Center for Molecular and Behavioral Neuroscience Rutgers University 197 University Ave, Newark, NJ 07120, USA.

http://dx.doi.org/10.1016/j.neuroimage.2016.11.037
Received 25 July 2016; Accepted 14 November 2016

cited study by Smith et al. (2011) aimed to address these concerns, by using a common generative model of the fMRI BOLD response (via the biophysically plausible Buxton-Friston balloon model; Buxton et al., 1998; Friston et al., 2000) to simulate a number of fMRI directed connectivity ground truths. The recovery of these ground truths was assessed across a variety of directional algorithms. Of these, the conditional Bayes method devised by Patel et al. (2006) (Patel's tau) was found to identify directed connections with the highest accuracy, albeit at an overall modest level (~65%), which raised broader questions as to the efficacy of applying directed connectivity to fMRI.

Later studies have re-examined the Smith simulations and provide a more optimistic view. For example, Ramsey et al. (2011), (2014) highlighted the Smith simulations' omission of algorithms that utilize multi-subject (i.e. group-level) data to identify directed connections, as well as the suppression of likely informative non-Gaussian signal components in their forward model and pre-processing steps. The authors hence applied their IMAGES (Independent Multiple-Sample Greedy Equivalence Search; Ramsey et al., 2010) group-level Bayes network algorithm, to the same Smith fMRI simulations, after removal of a high-pass Butterworth filter that actively suppressed non-Gaussian components in the original study. This yielded a marked improvement in directionality detection accuracy (> 85%; Ramsey et al., 2011, 2014). Deshpande and Hu (2012) highlighted further issues in the Smith forward model, which failed to include an explicit delay or lag in signaling between connected regions at the neuronal level. This might have contributed to the poor performance of lag-based Granger causality reported in the Smith simulations. Other fMRI simulations that included a realistic lag in neural signaling yielded far higher Granger detection accuracy (Roebroeck et al., 2005; Deshpande et al., 2010; Wang et al., 2014). The findings of these fMRI simulations should highlight general limitations of an overreliance on synthetic approaches to directed connectivity validation. Such simulations make generative assumptions that are open to debate, and which inevitably represent a simplification of the complexities of real imaging data.

Attempts to validate directed connectivity methods in MEG/EEG have been lacking in comparison to fMRI. This likely follows from the assumption that application of directed connectivity to MEG/EEG is fundamentally more apt, given their more direct measurement of neural activation (without complications arising from HRF convolution), sampled at a higher resolution and across a broader frequency spectrum. However, these temporal features come at the non-trivial cost of lower signal-to-noise ratio, increased non-stationarity and lower spatial resolution (i.e. the "inverse problem" of localizing neural sources for the raw sensor signal; Schoffelen and Gross, 2009). Clarifying how to address these unique challenges raises the need for directed connectivity validations in MEG/EEG as well as fMRI.

To this end, Wang et al., (2014) embedded directional ground truths in simulated MEG/EEG (via a neural mass model, Moran et al., 2013) and fMRI data (via the Buxton-Friston balloon model used in the Smith simulations), and compared the performance of a variety of directed connectivity algorithms across both modalities. The results support the efficacy of MEG/EEG directed connectivity analysis across a number of algorithms, as well as demonstrating comparably high detection performance in the fMRI simulations (again questioning the negative findings of the Smith simulations). Whilst these multi-modal simulations illustrate that convergent directed connectivity patterns are obtainable in MEG/EEG and fMRI data, clarification on more practical issues, such as improving signal-to-noise via pre-processing strategies and insight into appropriate significance testing are lacking. Similarly, the Wang et al. study failed to distinguish between sensor and source-level directed connectivity in their MEG/EEG forward model, and hence sidestepped the issue of "field spread" – the spreading of activity from a single neural source across proximal sensors, which has been shown to contaminate undirected connectivity analyses in MEG/EEG data (especially at the sensor-level; Schoffelen and Gross, 2009; Hipp et al., 2012).

Extending the undoubtedly useful fMRI and MEG/EEG simulation work calls for validations of directed connectivity in real data. Such validations are rare given the difficulty in specifying "empirical ground truth" directionality patterns in real compared to synthetic data. Whilst prior empirical directed connectivity studies have yielded interpretable results, both in fMRI (e.g. Mills-Finnerty et al., 2014; Wen et al., 2012, 2013) and in source-modeled MEG (e.g. Astolfi et al., 2007; Cole et al., 2010; Supp et al., 2007), these reports did not seek to address antecedent questions as to the base validity of applying directional methods to brain imaging data. Of the previous empirical validations, many focus on testing one specific algorithm, such as Granger causality (Roebroeck et al., 2005), Bayes network (Ramsey et al., 2014; Plis et al., 2011), directed coherence (Gómez-Herrero et al., 2008) or DCM (Bönstrup et al., 2016), rather than the more comprehensive multi-algorithm validations undertaken by the Smith and Wang simulations. Perhaps more problematic is that few if any of these prior validations have clearly formalized an a priori ground truth directed connectivity pattern with which to evaluate performance.

The over-reliance on simulations and the limited scope of the few prior empirical validations motivated the present report, which seeks to adapt the approach of the Smith and Wang simulations to affect a multi-algorithmic, multi-modal validation of directed connectivity in real data. We collected fMRI and (anatomically constrained, source-modeled) MEG data from the same sample of subjects, as they performed the same associative memory task, involving the cued retrieval of auditory-visual stimulus pairs (Fig. 1). The design of the task enabled testing of a common ground truth directed connectivity pattern that was predicated on widely replicated cognitive neuroscience findings and known patterns of anatomical connectivity. This capitalized on the established "sensory reactivation" effect in episodic memory research, wherein retrieval of auditory or visual information reactivates the same regions involved in perceiving those modalities (Slotnick and Schacter, 2004; Vaidya et al., 2002; Wheeler et al., 2000, 2006). By manipulating whether auditory stimuli cued retrieval of visual associates ("Aud-Vis" condition) or vice versa ("Vis-Aud" condition), we sought a ground truth reversal in directed connectivity between task conditions (i.e. auditory→visual ROI connectivity in the Aud-Vis condition, and visual→auditory ROI connectivity in the Vis-Aud condition; see Fig. 2). This ground truth is substantiated by known anatomical interconnectivity between auditory and visual regions identified in animals (Cappe and Barone, 2005; Mitani et al., 1985; Schroeder and Foxe, 2005).

The ground truth was tested in algorithms spanning a diverse set of assumptions, without any bias or emphasis placed on validating a given algorithm (or imaging modality). The tested algorithms can broadly be categorized as "pairwise", if they orient on an isolated connection-by-connection basis (Granger causality, Patel's tau and phase slope index), or "multivariate", if they orient individual connections only after considering all other conditioning relationships amongst the region time series (IMAGES; see Supplementary Materials for descriptions of all tested algorithms).[1] Note that these selected algorithms span a core rather than exhaustive set, and were chosen on the basis of being readily available (i.e. implementable) and previously applied to human imaging data, whilst still inferring directionality via fairly distinct mathematical assumptions. Whilst other directional algorithms that have performed well in prior simulations could also have been included, we chose to only test these core algorithms so as to fulfill both our aims – not only to validate the use of directed connectivity algorithms in human imaging data, but also to highlight the virtue of the general approach to validating via empirical rather than simulated "ground truths". Emphasizing the latter point might encourage future

---

[1] Note that, given recent emphasis on large network graphs for characterizing brain "connectomes" (Sporns, 2014), we focus on directional algorithms that are not limited to modeling a small number of regions (unlike, e.g., DCM).