Journal of Neuroscience Methods 268 (2016) 142–153

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

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

Multivariate dynamical systems-based estimation of causal brain interactions in fMRI: Group-level validation using benchmark data, neurophysiological models and human connectome project data

Srikanth Ryali[a,*], Tianwen Chen[a], Kaustubh Supekar[a], Tao Tu[a], John Kochalka[a], Weidong Cai[a], Vinod Menon[a,b,c,*]

[a] Department of Psychiatry & Behavioral Sciences, Stanford University School of Medicine, Stanford, CA 94305, United States
[b] Department of Neurology & Neurological Sciences, Stanford University School of Medicine, Stanford, CA 94305, United States
[c] Stanford Neurosciences Institute, Stanford University School of Medicine, Stanford, CA 94305, United States

HIGHLIGHTS

- Validated MDS using an unbiased approach with simulation models different from the estimation models.
- Validation datasets consists of benchmark as well as datasets simulated from a stochastic neurophysiological model.
- Examined the stability of causal interactions in a fronto-cingulate-parietal control network in a working memory task.
- MDS is effective in estimating causal interactions in both the simulation datasets.
- The stability analysis revealed that the right insula functions as a causal hub during working memory.

ARTICLE INFO

Article history:
Received 25 November 2015
Received in revised form 11 March 2016
Accepted 13 March 2016
Available online 22 March 2016

Keywords:
BOLD fMRI
Causality
DCM
GCA

ABSTRACT

Background: Causal estimation methods are increasingly being used to investigate functional brain networks in fMRI, but there are continuing concerns about the validity of these methods.
New method: Multivariate dynamical systems (MDS) is a state-space method for estimating dynamic causal interactions in fMRI data. Here we validate MDS using benchmark simulations as well as simulations from a more realistic stochastic neurophysiological model. Finally, we applied MDS to investigate dynamic casual interactions in a fronto-cingulate-parietal control network using human connectome project (HCP) data acquired during performance of a working memory task. Crucially, since the ground truth in experimental data is unknown, we conducted novel stability analysis to determine robust causal interactions within this network.
Results: MDS accurately recovered dynamic causal interactions with an area under receiver operating characteristic (AUC) above 0.7 for benchmark datasets and AUC above 0.9 for datasets generated using the neurophysiological model. In experimental fMRI data, bootstrap procedures revealed a stable pattern of causal influences from the anterior insula to other nodes of the fronto-cingulate-parietal network.
Comparison with existing methods: MDS is effective in estimating dynamic causal interactions in both the benchmark and neurophysiological model based datasets in terms of AUC, sensitivity and false positive rates.
Conclusions: Our findings demonstrate that MDS can accurately estimate causal interactions in fMRI data. Neurophysiological models and stability analysis provide a general framework for validating computational methods designed to estimate causal interactions in fMRI. The right anterior insula functions as a causal hub during working memory.

* Corresponding authors at: Department of Psychiatry & Behavioral Sciences, 401 Quarry Rd., Stanford University School of Medicine, Stanford, CA 94305-5719, United States. Fax: +1 650 736 7200.
E-mail addresses: sryali@stanford.edu (S. Ryali), menon@stanford.edu (V. Menon).

http://dx.doi.org/10.1016/j.jneumeth.2016.03.010
0165-0270/

1. Introduction

There is a growing interest in examining how cognitive functions emerge as a result of dynamic causal interactions among distributed brain regions. Computational methods for estimating dynamic causal interactions from fMRI data include state-space dynamical models (Daunizeau et al., 2009; Friston et al., 2003; Li et al., 2011; Ryali et al., 2011; Smith et al., 2009), Granger causal analysis (GCA) (Barnett and Seth, 2013; Deshpande et al., 2008; Goebel et al., 2003; Jiao et al., 2011; Roebroeck et al., 2005; Seth, 2010; Wen et al., 2012), structural equation modeling (SEM) (Gates and Molenaar, 2012; Gates et al., 2010; Gates et al., 2011; McIntosh and Gonzalez-Lima, 1994) and Bayesian network methods (Ramsey et al., 2011; Ramsey et al., 2009). Despite much progress in the field, there is a growing debate about the validity of these methods. In an attempt to address this issue, Smith et al. (2011) evaluated the performance of several methods, including GCA and Bayesian network methods, on several simulated datasets at a single-subject level. The performance of these methods was, however, not assessed at the group level. This remains a critical gap in the literature as almost all human fMRI studies are based on inferences on data from multiple participants (Gates and Molenaar, 2012; Ramsey et al., 2011; Schippers et al., 2011).

The performance of state-space based causal estimation methods (Cai et al., 2015; Chen et al., 2014; Ryali et al., 2011; Supekar and Menon, 2012) that overcome limitations of existing methods has not been validated using benchmark datasets or other more biologically plausible datasets generated independent of the methods used to test them. Critically, estimating causal interactions in fMRI is challenging because (a) neuronal interactions occur in the range of 20–50 ms while fMRI signals are sampled at 2–3 s, and (b) different brain regions have varying hemodynamic response (HRFs) that link neuronal signals to the observed fMRI response (Friston et al., 2003; Ryali et al., 2011; Seth et al., 2013; Smith et al., 2009). These factors represent a major challenge for computational methods designed to infer causal interactions in experimental fMRI data. Here we use extensive simulations and stability analysis to evaluate the performance of our previously developed multivariate dynamical systems (MDS) state-space methods (Ryali et al., 2011) at the group level on three different types of datasets: (1) benchmark data provided by Smith et al. (2011), (2) fMRI datasets simulated using a stochastic neurophysiological model (Holcman and Tsodyks, 2006; Testa-Silva et al., 2012) that improves upon the models implemented by Smith et al. (2011), and (3) experimental fMRI data on a working memory task from a group of 63 participants, acquired as part of the human connectome project (HCP).

MDS is a state-space approach for estimating dynamic causal interactions in fMRI data at both single subject and group level. MDS estimates event-specific causal interactions in fMRI while accounting for regional variations in HRF across brain areas and individuals (Ryali et al., 2011). Previously, we evaluated the performance of MDS using data as a function of signal to noise ratio (SNR), regional variation in HRF characteristics, network size, number of observations and experimental design (Ryali et al., 2011). We found that MDS accurately estimates causal interactions and that its performance is much better than GCA. However, our use of MDS-based models for generating simulated test datasets could potentially have biased the evaluation of its performance. This is a common problem in many studies seeking to validate new computational methods in the neuroimaging literature. For example, in a recent analysis of group level GCA, model performance was evaluated using a very simplistic two-node network (Schippers et al., 2011). The simulated datasets were generated using a bivariate autoregressive (AR) model which in turn was used for estimating the Granger causality between the two nodes. It is therefore not clear how group level GCA performs when the simulation model is different from the model used for estimating the causal interactions; for example, when a more realistic neuronal mass model is used to stimulate the fMRI data. To avoid potential biases when evaluating the performance of individual methods, it is critical to use data that is generated independently from the models used to estimate dynamical causal interactions (Seth et al., 2013; Smith et al., 2011).

Here we address three critical issues associated with the validation of MDS: (1) unbiased evaluation using simulated data generated from models different from MDS, (2) simulation of more realistic datasets with causal interactions occurring at a 20–50 ms time resolution at the neuronal level with BOLD-fMRI signals downsampled to 2–3 s, and (3) validation of MDS at the group level, similar to experimental fMRI studies. We used datasets generated using two simulation models: (i) previously published benchmark datasets generated using a deterministic bi-linear DCM model (Friston et al., 2003; Smith et al., 2011) and (ii) new datasets generated using a more biologically realistic neurophysiological model that incorporates nonlinear interactions between neuronal populations, conduction delays and saturation effects (Holcman and Tsodyks, 2006; Testa-Silva et al., 2012). In the latter case, a reduced stochastic dynamical system is used to model an interconnected network of excitatory neurons with activity-dependent synaptic depression (Holcman and Tsodyks, 2006). Crucially, this model system reproduces key aspects of intrinsic cortical dynamics, including spontaneous state transitions, and is well suited to examining brain dynamics in experimental preparations with both typical and atypical synaptic connections (Holcman and Tsodyks, 2006; Testa-Silva et al., 2012). Here we extend the original neurophysiological model by creating 5, 10 and 15 node networks and incorporating biologically realistic delays into inter-node signaling to generate casual interactions within a stochastic dynamical system framework. We show through extensive simulations, using both deterministic DCM and stochastic neurophysiological models, that MDS can accurately recover dynamic causal networks in simulated fMRI data.

Despite their advantages, in principle, simulations cannot model all aspects of fMRI data. To address this issue and further validate MDS, we applied MDS on fMRI data acquired during performance of a working memory task and examined causal interactions within a fronto-cingulate-parietal network important for cognitive control (Cai et al., 2015; Chen et al., 2014; Dosenbach et al., 2008; Dosenbach et al., 2007; Menon, 2011; Menon and Uddin, 2010; Supekar and Menon, 2012). This control network includes anterior insula (AI), anterior cingulate cortex (ACC), ventrolateral prefrontal cortex (VLPFC), dorsolateral prefrontal cortex (DLPFC) and posterior parietal cortex (PPC) regions that are co-activated across a wide range of cognitive tasks (Menon, 2011; Menon and Uddin, 2010). Critically, the AI node of this network is thought to play a key role in switching between large-scale brain networks and facilitating access to attention and working memory resources (Sridharan et al., 2007; Sridharan et al., 2008). Previous studies using a variety of different computational methods have reported that the AI has a dominant causal influence on other prefrontal, cingulate and parietal regions during tasks involving orienting attention and response inhibition (Cai et al., 2015; Chen et al., 2014; Ham et al., 2013; Sridharan et al., 2008; Supekar and Menon, 2012). Here, we use open-source task fMRI data (Van Essen et al., 2012) acquired during a working memory task which required participants to continually encode, maintain and update information in mind (Baddeley, 1996). Based on the research reviewed above, we predicted that the AI would play a dominant causal role during working memory task performance with significant influence on other nodes of the fronto-cingulate-parietal network. Crucially, because the ground truth in experimental data is not known, we used novel stability analysis of data from 63 participants to identify robust causal influences within this network. In sum, we demonstrate that MDS can