NeuroImage 161 (2017) 134–148

Contents lists available at ScienceDirect

NeuroImage

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

A Bayesian spatial model for neuroimaging data based on biologically informed basis functions

Ismael Huertas [a], Marianne Oldehinkel [b,c], Erik S.B. van Oort [c], David Garcia-Solis [d,e], Pablo Mir [a,e], Christian F. Beckmann [b,c,f,1], Andre F. Marquand [b,c,g,*,1]

[a] Unidad de Trastornos del Movimiento, Servicio de Neurología y NeurofisiologíaClínica, Instituto de Biomedicina de Sevilla (IBiS), Hospital UniversitarioVirgen del Rocío/CSIC/Universidad de Sevilla, Seville, Spain
[b] Department of Cognitive Neuroscience, Radboud University Medical Centre, Nijmegen, The Netherlands
[c] Donders Centre for Cognitive Neuroimaging, Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, The Netherlands
[d] Servicio de Medicina Nuclear, UDIM, Hospital UniversitarioVirgen del Rocío, Seville, Spain
[e] Centro de Investigación Biomédica en Red sobre Enfermedades Neurodegenerativas (CIBERNED), Spain
[f] Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB), University of Oxford, United Kingdom
[g] Department of Neuroimaging, Centre for Neuroimaging Sciences, Institute of Psychiatry, King's College London, United Kingdom

ARTICLE INFO

Keywords:
Multivariate GLM
Functional parcellations
Spatial statistics
Basis functions
Spatial statistics
Dopamine transporter SPECT
Parkinsonian disorders

ABSTRACT

The dominant approach to neuroimaging data analysis employs the voxel as the unit of computation. While convenient, voxels lack biological meaning and their size is arbitrarily determined by the resolution of the image. Here, we propose a multivariate spatial model in which neuroimaging data are characterised as a linearly weighted combination of multiscale basis functions which map onto underlying brain nuclei or networks or nuclei. In this model, the elementary building blocks are derived to reflect the functional anatomy of the brain during the resting state. This model is estimated using a Bayesian framework which accurately quantifies uncertainty and automatically finds the most accurate and parsimonious combination of basis functions describing the data. We demonstrate the utility of this framework by predicting quantitative SPECT images of striatal dopamine function and we compare a variety of basis sets including generic isotropic functions, anatomical representations of the striatum derived from structural MRI, and two different soft functional parcellations of the striatum derived from resting-state fMRI (rfMRI). We found that a combination of ~50 multiscale functional basis functions accurately represented the striatal dopamine activity, and that functional basis functions derived from an advanced parcellation technique known as Instantaneous Connectivity Parcellation (ICP) provided the most parsimonious models of dopamine function. Importantly, functional basis functions derived from resting fMRI were more accurate than both structural and generic basis sets in representing dopamine function in the striatum for a fixed model order. We demonstrate the translational validity of our framework by constructing classification models for discriminating parkinsonian disorders and their subtypes. Here, we show that ICP approach is the only basis set that performs well across all comparisons and performs better overall than the classical voxel-based approach. This spatial model constitutes an elegant alternative to voxel-based approaches in neuroimaging studies; not only are their atoms biologically informed, they are also adaptive to high resolutions, represent high dimensions efficiently, and capture long-range spatial dependencies, which are important and challenging objectives for neuroimaging data.

1. Introduction

Neuroimaging techniques have become invaluable tools for clinical research and practice in many brain disorders thanks to their ability to noninvasively investigate brain structure and function with relatively high spatial resolution. Data acquisition techniques such as MRI and PET allow the rich spatial structure that emerges from interactions between brain regions to be probed in high detail. However, the predominant analysis approaches that rely on the voxel as the unit of analysis do not take full advantage of this source of information. In the classical mass-

* Corresponding author. Donders Centre for Cognitive Neuroimaging, Donders Institute for Brain, Cognition and Behaviour, Kapittelweg 29, 6525 EN, Nijmegen, The Netherlands.
E-mail address: a.f.marquand@fcdonders.ru.nl (A.F. Marquand).
[1] Shared senior author.

http://dx.doi.org/10.1016/j.neuroimage.2017.08.009
Received 22 June 2017; Received in revised form 27 July 2017; Accepted 2 August 2017
Available online 4 August 2017

univariate approach, which entails fitting independent temporal models at each sampled brain location (i.e. each voxel), spatial dependencies are effectively disregarded or dealt with suboptimally (e.g. by smoothing the data). This ignores an important source of information encoded by statistical dependencies between brain regions. The mass-univariate approach also generates a large number of statistical estimates that depend arbitrarily on the voxel size in the image. These spatially uninformed estimates need to be combined and inferred upon using complex post-hoc correction methods such as random field theory (Nichols, 2012; Worsley et al., 1996), the accuracy of which has been recently called into question (Eklund et al., 2016). Voxel-based features are also potentially suboptimal for multivariate approaches such as pattern recognition (Wolfers et al., 2015; Mwangi et al., 2014) essentially because voxels lack biological meaning. While pattern recognition approaches can make use of correlations between brain regions, the nature of neuroimaging data often leads to severely ill-posed problems (e.g. with hundreds of thousands of features and tens to hundreds of samples). Therefore, whole-brain voxel-based approaches are not optimal for discriminating conditions if the underlying signal is localized to particular regions or networks. For multivariate approaches as well as mass-univariate approaches it is therefore desirable to find parsimonious representations of brain structure or function that can more faithfully represent the underlying signal. Such models may predict clinically-relevant outcomes more accurately than voxel-based approaches and may be more interpretable in the sense that discriminating features may be cleanly related to underlying neuronal units of computation.

In light of these considerations, there have been some proposals to take spatial dependencies into account using multivariate approaches, and the field of spatial statistics offers attractive methods in this respect. Various discrete spatial models have been proposed for neuroimaging data (e.g., Penny et al., 2005; Woolrich et al., 2004) but these generally only provide local smoothing for the parameter estimates from mass-univariate analysis. They do not accommodate long-range dependencies that are intrinsic to neuroimaging data, nor overcome the arbitrary dependence on voxel size or the intricate structure-shape relationships of the brain. A more accurate and flexible approach is the spatial mixed model, in which an additional term, called a spatial random effect, is added to the model. Here, spatial dependencies are typically modeled using a continuous (usually Gaussian) spatial random field. The covariance matrix of this term describes the spatial correlation between allocations (e.g., voxels), and the inversion of this matrix is necessary to obtain suitable estimates under this model (Wikle and Royle, 2002). The immediate problem of applying this approach to neuroimaging data is the computational burden of this matrix inversion. Accordingly, this approach has principally been used in the context of restricted regions of interest (Bowman et al., 2008; Groves et al., 2009) although some studies have made use of data reduction techniques to approximate the underlying spatial process (Hyun et al., 2014; Zhu et al., 2014). An efficient alternative to model high-dimensional spatial processes is the use of low rank models, in which the covariance matrix is approximated by a reduced number of basis functions (Cressie and Johannesson, 2008). Most commonly, these basis functions are taken to be nonlinear functions, such as radial basis functions (RBFs), b-splines, or wavelets, that are placed all over the spatial domain. In spatial applications, multiple resolutions are typically used to capture both short and long ranges of spatial dependencies.

In this work, we introduce a spatial statistical modelling framework that uses data-driven basis functions to model neuroimaging data. These basis functions are derived from measures of brain function, and therefore more closely reflect the underlying biology relative to generic basis functions. While various spatial basis sets are possible, we propose to use a soft multiscale functional parcellation derived from resting-state fMRI (rfMRI). For this, we employ a parcellation strategy known as Instantaneous Connectivity Parcellation (ICP, van Oort et al., 2016). Our rationale is based on emerging evidence of temporally independent, spatially overlapping, subnetworks within anatomical regions and functional networks in the human brain (Smith et al., 2012). These subnetworks are believed to represent fine-scale units of computation used by the brain for processing. We use these subnetworks as basis function because of their correspondence with biology. There are various strategies that we could employ to extract these subnetworks (e.g., Yeo et al., 2011; Craddock et al., 2012; Shen et al., 2013; Gordon et al., 2016; Glasser et al., 2016), but the ICP approach is well suited to deriving such subnetworks as it combines three features: first, ICP sub-divides brain networks on the basis of fine-grained temporal similarities instead of temporally averaged correlations. Second, ICP does not impose a spatial contiguity constraint, meaning that brain regions that are not spatially adjacent can still participate in the same subnetwork. Finally, ICP follows a top-down strategy for parcellation, which generates sets of parcels at different levels of granularity which allows us to model multiple ranges of spatial dependencies in the image. We compare this approach to a variety of basis sets including: i) generic isotropic bisquare functions commonly used in spatial applications (Cressie and Johannesson, 2008); ii) structural parcellations of the striatum derived from two different atlases; and iii) functional parcellations of the striatum obtained from Independent Component Analysis (ICA).

For model fitting, we propose to use a Bayesian regression framework to automatically find a linearly weighted sum of basis functions that accurately fits an imaged brain region (or to the whole brain). The resulting basis function fit and the corresponding weights can be used in a second level of analysis to investigate the phenotype of the imaged subjects. To illustrate, we test our framework to predict quantitative SPECT data of the dopamine transporter (DAT) availability in the healthy striatum. DAT imaging allows assessing the integrity of presynaptic dopaminergic neurons of the nigrostriatal pathway and it is widely used in the clinical practice of movement disorders (Tatsch and Poepperl, 2013). We provide an example of how this method can be applied to a real clinical application. For this, we use the DAT data to automatically differentiate between different diagnosed sub-cohorts corresponding to different parkinsonian disorders. We hypothesized that spatial models that are informed by brain function would be superior to spatial models that are informed only by the structural anatomy and to generic models that do not incorporate knowledge of the underlying biology. Therefore, we compare functionally informed basis functions derived from resting state fMRI to anatomical basis functions derived from two widely used anatomical parcellations of the striatum and also to generic basis functions commonly used in spatial applications. The clinical application we have chosen provides an exacting test of this hypothesis for three reasons: (i) the spatial resolution of SPECT is low relative to alternative methods (e.g. fMRI) meaning that clinically relevant spatial dependencies are difficult to detect; (ii) anatomical subdivisions are well-defined for the striatum, which biases the analysis in favour of anatomical parcellations and (iii) the data modality used to create the basis set (BOLD fMRI, indirectly measuring oxygen consumption) measures different aspects of the underlying biology relative to the clinical biomarker (DAT SPECT, measuring dopamine function). Therefore the method must learn dependencies that generalize across different aspects of brain function.

Our approach is related to several lines of work in the neuroimaging literature. Gershman et al. (2011) developed a spatial modeling approach for neuroimaging data, referred to as topographic latent source analysis (TLSA). In TLSA, fMRI data are modeled as a superposition of image sources constructed from adaptive RBFs. Like our approach, TLSA abstracts away from the voxel as a unit of analysis, instead performing inferences over underlying neuroanatomical regions. However, in TLSA generic isotropic RBFs are used that do not map cleanly onto their biological sources (i.e. brain nuclei). The approach also requires running heavy optimization machinery in order to fit a given data set. Our approach is also related to dictionary learning approaches (e.g. Varoquaux et al., 2011) and to approaches that model neuroimaging data using multi-scale parcellations (e.g. Jennaton et al., 2012; Bellec, 2013). These approaches generally aim to segment a set of neuroimaging data into subject-specific or group level atlases. In contrast, our approach