NeuroImage 112 (2015) 232–243

Contents lists available at ScienceDirect

NeuroImage

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

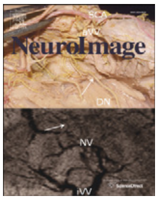

Gaussian process classification of Alzheimer's disease and mild cognitive impairment from resting-state fMRI

Edward Challis [a], Peter Hurley [a], Laura Serra [c], Marco Bozzali [c], Seb Oliver [a], Mara Cercignani [b],*

[a] *Department of Physics and Astronomy, University of Sussex, Falmer, East Sussex BN1 9QH, UK*
[b] *Clinical Imaging Sciences Centre, Brighton and Sussex Medical School, University of Sussex, Falmer, East Sussex BN1 9PR, UK*
[c] *Neuroimaging Laboratory, Santa Lucia Foundation, Via Ardeatina 306, Roma, Italy*

ARTICLE INFO

Article history:
Accepted 17 February 2015
Available online 28 February 2015

Keywords:
Machine learning
Functional connectivity
Dementia

ABSTRACT

Multivariate pattern analysis and statistical machine learning techniques are attracting increasing interest from the neuroimaging community. Researchers and clinicians are also increasingly interested in the study of functional-connectivity patterns of brains at rest and how these relations might change in conditions like Alzheimer's disease or clinical depression. In this study we investigate the efficacy of a specific multivariate statistical machine learning technique to perform patient stratification from functional-connectivity patterns of brains at rest. Whilst the majority of previous approaches to this problem have employed support vector machines (SVMs) we investigate the performance of Bayesian Gaussian process logistic regression (GP-LR) models with linear and non-linear covariance functions. GP-LR models can be interpreted as a Bayesian probabilistic analogue to kernel SVM classifiers. However, GP-LR methods confer a number of benefits over kernel SVMs. Whilst SVMs only return a binary class label prediction, GP-LR, being a probabilistic model, provides a principled estimate of the probability of class membership. Class probability estimates are a measure of the confidence the model has in its predictions, such a confidence score may be extremely useful in the clinical setting. Additionally, if miss-classification costs are not symmetric, thresholds can be set to achieve either strong specificity or sensitivity scores. Since GP-LR models are Bayesian, computationally expensive cross-validation hyper-parameter grid-search methods can be avoided. We apply these methods to a sample of 77 subjects; 27 with a diagnosis of probable AD, 50 with a diagnosis of a-MCI and a control sample of 39. All subjects underwent a MRI examination at 3 T to obtain a 7 minute and 20 second resting state scan. Our results support the hypothesis that GP-LR models can be effective at performing patient stratification: the implemented model achieves 75% accuracy disambiguating healthy subjects from subjects with amnesic mild cognitive impairment and 97% accuracy disambiguating amnesic mild cognitive impairment subjects from those with Alzheimer's disease, accuracies are estimated using a held-out test set. Both results are significant at the 1% level.

Introduction

A broad goal of neuroimaging research is to develop effective, reliable clinical tools for the early detection and diagnosis of a range of neurological conditions such as dementia, depression and attention deficit hyperactivity disorder. Machine learning (ML) seems a promising route to help achieve such objectives. ML is the study of algorithms and computational techniques that use previous examples in the form of multivariate datasets to help make future predictions. One application of a ML prediction algorithm in the context of neuroimaging would be to make clinical diagnoses from subject's functional MRI (fMRI) scans. Alongside providing a computational and statistical framework within which to make predictions from multivariate observations, ML can also provide insights into what multivariate features of the data are most relevant for making accurate predictions. In the context of neuroimaging for patient stratification those features correspond to biomarkers of disease states.

In this paper we present a ML technique to perform patient stratification between healthy control subjects and either amnesic mild cognitive impairment (a-MCI) or Alzheimer's disease subjects. Subject classifications are made from the functional-connectivity scores of their brains inferred from resting state fMRI (rsfMRI) scans. Previous rsfMRI patient stratification studies have applied support vector machines (SVMs) to make inter-group classifications. Our approach here

---

Abbreviations: AC, accuracy; AD, Alzheimers Disease; a-MCI, amnestic mild cognitive impairment; ARD, automatic relevance determination; AUC, area under the ROC curve; DMN, default mode network; GP-LR, Gaussian process logistic regression; ICA, independent component analysis; LOOCV, leave one out cross validation; MMSE, mini-mental state examination; NC, normal control; PCA, principal component analysis; SC, specificity; SS, sensitivity; SVM, support vector machine

* Corresponding author at: Clinical Imaging Sciences Centre, University of Sussex, Falmer, Brighton BN1 9RR, UK. Fax: +44 1273 876721.

E-mail address: m.cercignani@bsms.ac.uk (M. Cercignani).

http://dx.doi.org/10.1016/j.neuroimage.2015.02.037
1053-8119/

is to use Gaussian process logistic regression (GP-LR) models. Whilst examples of application of Gaussian processes to neuroimaging data exist (e.g. Marquand et al., 2010), we are unaware of any other functional-connectivity studies using Gaussian process models to make such group level predictions.

The motivations of our work are two-fold. First, to show how GP-LR models can be applied to inter-group rsfMRI classification problems and to highlight some advantages of this approach. Second, to investigate what classification accuracy such a technique can achieve in distinguishing between healthy, a-MCI and AD subjects, and to identify what features of the data are most relevant in driving those predictions.

Machine learning fMRI studies

A neuroimaging problem to which ML can be applied is to predict whether a subject belongs to one of a number of different subject groups, for example to predict if a subject is healthy versus diseased or young versus old. Due to inter-subject, inter-scan and inter-centre variability, and the often limited number of example scans that are available to researchers, this is typically a hard statistical inference problem. In this study we seek to address the inter-group prediction problem. We apply ML methods with the specific aim to automatically disambiguate healthy control subjects from subjects exhibiting symptoms of amnesic mild cognitive impairment (a-MCI) and Alzheimer's disease (AD). Importantly, our data is not longitudinal; each scan corresponds to a different individual.

A good introduction to machine learning methods applied to neuroimaging problems can be found in the review articles by Pereira et al. (2009), Lemm et al. (2011) and Ashburner and Klöppel (2011). A more general introduction to probabilistic machine learning and Bayesian methods can be found in Barber (2012).

Resting-state functional MRI

We seek to perform patient stratification from the application of ML algorithms to resting-state fMRI scans. RsfMRI data refers to fMRI scans that are recorded whilst the subject is at rest; that is, the subject is not performing any particular task and is not asleep. From a practical perspective, resting-state scans have the advantage of being easier to acquire than scans recorded whilst the subject is performing a task because fewer experimental variables have to be controlled for. Thus, inter-scan differences that are not attributable to the subjects' mental state are minimised and group differences will be easier to infer. Furthermore, since many subjects, such as those that have Alzheimer's disease or dementia, are often incapable of carrying out cognitive tasks required by task-based studies, resting-state studies have the benefit of being able to include such subjects without biasing the experimental design.

RsfMRI voxel blood-oxygenation-level-dependent (BOLD) signal time-courses exhibit low frequency, ($\approx$0.1 Hz), oscillations. These spontaneous BOLD signal oscillations exhibit temporal correlations across spatially distinct brain regions. Such patterns of activity are now commonly believed to mirror the functional-connectivity patterns of the brain (Van Den Heuvel and Hulshoff Pol, 2010). Assuming that these patterns reflect specific resting state networks, one of them, namely the default mode network (DMN) has received particular attention. Evidence suggests that during goal directed behaviour the DMN correlations are suppressed (Buckner et al., 2008; Gusnard and Raichle, 2001). Multiple studies have observed that changes to the DMN may be biomarkers for various neurological conditions such as Alzheimer's disease (Koch et al., 2012; Greicius et al., 2004b), attention deficit hyperactivity disorder (Uddin et al., 2008; Liddle et al., 2011) and depression (Zeng et al., 2012; Sheline et al., 2009; Bluhm et al., 2009) amongst other studies.

Functional-connectivity

Functional-connectivity is commonly defined as the temporal dependence of neuronal activity patterns of anatomically separated brain regions (Friston et al., 1993). As such, functional-connectivity is a property of the brain that is static and independent of time. Whole brain resting-state functional-connectivity patterns are obtained by studying the coactivation between the time-courses of voxels, or collections of voxels, that are spatially distributed. Typically, the methods that are employed to discern functional-connectivity relations from rsfMRI data fall into two categories: Model-free methods such as independent component analysis or principal component analysis and model-based methods such as region of interest or seed correlation analysis. See Cole et al. (2010) for an introductory review of these techniques.

Model-free methods aim to find a reduced set of temporal basis functions such that each voxel's BOLD time-course can be well approximated by a liner combination of these temporal bases. The temporal basis functions are most frequently estimated using either the principal component analysis (PCA) or the independent component analysis (ICA) statistical models. Having applied PCA or ICA, functional-connectivity between two anatomically distinct regions is inferred if the two groups share similar temporal basis function coefficients. ICA and PCA methods are thought of as model-free in the sense that no brain region atlas is defined by the researcher a priori. However, the temporal bases are found by fitting a statistical model which makes certain assumptions about the data generating process, for example PCA finds the basis functions that span the directions of maximum variance and ICA finds basis functions that span the directions that maximise the kurtosis (or some other proxy of statistical independence). In this sense the model-free label is misleading. A practical consideration when using ICA or PCA methods is that the temporal bases can be difficult to interpret — deciding whether a basis is due to 'noise' or neuronal variability is typically decided by a human expert. Automatically ordering and labelling the temporal bases is the subject of on-going methodological research (Tohka et al., 2008; De Martino et al., 2007). A further difficulty with applying model-free methods as a data preprocessing step for making inter-group predictions is that it is unclear whether the temporal bases calculated from one group generalise to another. These issues make it difficult to apply model-free methods as a functional-connectivity preprocessing step in a ML system designed to make inter-group predictions.

An alternative to model-free methods are so called model-based methods. Model-based methods infer functional-connectivity by inspecting the temporal dependence in BOLD signals between anatomically distinct brain regions. Whilst many different time-course dependence metrics could be used to infer functional-connectivity (Zhou et al., 2009), a commonly used and simple metric is the spontaneous correlation in BOLD signals between brain regions. In such an analysis regions that have highly correlated time-courses are inferred to be functionally connected. We refer to this approach as the regions of interest (ROIs) method. Other names used in the literature include volumes of interest or seed based correlation analysis. These techniques are thought of as model-based because the seed ROIs need to be specified a priori and so connectivity patterns are not directly inferred from the data. The primary strength of this approach is the ease with which it can be implemented and the results interpreted. Thus, the ROI approach is the favoured functional-connectivity preprocessing technique for patient stratification ML studies (Craddock et al., 2009; Zeng et al., 2012; Meier et al., 2012; Anderson et al., 2011). There is also some evidence, specific to the problem of disambiguating healthy versus AD subjects, that model-based methods may have more diagnostic power than model-free methods (Koch et al., 2012). In this work, the authors hypothesise that model-based methods may have more diagnostic power due to correlational analysis being more robust to BOLD signal variability that is observed to increase with age or the partial volume effects of grey matter loss.