NeuroImage xxx (2016) xxx–xxx

Contents lists available at ScienceDirect

NeuroImage

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

Multivariate and repeated measures (MRM): A new toolbox for dependent and multimodal group-level neuroimaging data

Martyn McFarquhar [a,*], Shane McKie [a], Richard Emsley [b], John Suckling [c], Rebecca Elliott [a], Stephen Williams [d]

[a] Neuroscience & Psychiatry Unit, Stopford Building, The University of Manchester, Oxford Road, Manchester M13 9PL, UK
[b] Centre for Biostatistics, Jean McFarlane Building, The University of Manchester, Oxford Road, Manchester M13 9PL, UK
[c] Brain Mapping Unit, Herchel Smith Building for Brain and Mind Sciences, University of Cambridge, Robinson Way, Cambridge CB2 0SZ, UK
[d] Imaging Sciences, Stopford Building, The University of Manchester, Oxford Road, Manchester M13 9PL, UK

ARTICLE INFO

Article history:
Received 10 December 2015
Accepted 17 February 2016
Available online xxxx

Keywords:
Multivariate GLM
Permutation
Multimodal
Repeated measures
Discriminant functions

ABSTRACT

Repeated measurements and multimodal data are common in neuroimaging research. Despite this, conventional approaches to group level analysis ignore these repeated measurements in favour of multiple between-subject models using contrasts of interest. This approach has a number of drawbacks as certain designs and comparisons of interest are either not possible or complex to implement. Unfortunately, even when attempting to analyse group level data within a repeated-measures framework, the methods implemented in popular software packages make potentially unrealistic assumptions about the covariance structure across the brain. In this paper, we describe how this issue can be addressed in a simple and efficient manner using the multivariate form of the familiar general linear model (GLM), as implemented in a new MATLAB toolbox. This multivariate framework is discussed, paying particular attention to methods of inference by permutation. Comparisons with existing approaches and software packages for dependent group-level neuroimaging data are made. We also demonstrate how this method is easily adapted for dependency at the group level when multiple modalities of imaging are collected from the same individuals. Follow-up of these multimodal models using linear discriminant functions (LDA) is also discussed, with applications to future studies wishing to integrate multiple scanning techniques into investigating populations of interest.

1. Introduction

Group-level repeated measurements are commonplace in neuroimaging research, from neurocognitive paradigms with multiple activation conditions to longitudinal intervention studies. Despite this, conventional summary statistic approaches to modelling these data ignore the repeated measurements in favour of the construction of contrasts at the subject level. These contrasts are then explored using multiple group-level linear models. Though this approach is advantageous due to its simplicity, when the design contains more than two repeated measurements many of the typical ANOVA tests used to investigate the repeated measures and their interactions are either overly complex to implement or simply not possible. Furthermore, for approaches such as the *p*-block method of analysing pharmacological challenge fMRI data (phMRI; e.g. McKie et al., 2011), the use of contrasts at the individual-level is not a useful method and repeated-measurement models become a necessity. Despite this, the approaches currently implemented in two of the most popular fMRI analysis packages, FSL (http://fsl.fmrib.ox.ac.uk/fsl/) and SPM (http://www.fil.ion.ucl.ac.uk/spm/), are not able to easily account for dependent group-level neuroimaging data. FSL FEAT must assume sphericity at every voxel so that *F*-tests follow an exact *F*-distribution (Huynh and Feldt, 1970). Cases where the sphericity condition is not met can lead to a poorer control of the type I error rate due to overly liberal *F*-statistics (Box, 1954; Kogan, 1948). SPM, on the other hand, has a method for correcting departures from sphericity (Glaser and Friston, 2007). However, the estimated structure used in this correction is assumed to be the same for every voxel. In both cases, these assumptions may not always be valid for complex dependent data.

Further to the issues of dependent group-level analyses, it is also commonplace to collect multiple imaging sequences from the same subjects during the same scanning session (e.g. functional, T1 structural, arterial spin labelling). In some cases, there may even be different modalities of imaging collected from the same individuals (e.g. MR and PET). Analysing these different sequences/modalities is similar to repeated-measures designs due to the assumed correlation between measurements taken from the same individual. The biggest difference with repeated-measurement models is simply that the data are not guaranteed to be commensurate as they are generally not measured on the same scale. Although questions of interest often focus on the sequences and modalities individually, pooling the information provided by different imaging techniques may be advantageous in exploring

* Corresponding author at: Neuroscience and Psychiatry Unit, University of Manchester, G.708 Stopford Building, Oxford Road, Manchester M13 9PL, UK. Tel: +44 161 275 5775.
E-mail address: martyn.mcfarquhar@manchester.ac.uk (M. McFarquhar).

http://dx.doi.org/10.1016/j.neuroimage.2016.02.053
1053-8119/

how a combination of measurements may provide information on group differences above and beyond the information they provide individually. To achieve this, methods that accommodate both the assumed correlation and the differing scales of the measurements are needed.

In this paper, we will demonstrate how both the issues of repeated-measures and multimodal[1] group models can be addressed using the multivariate form of the familiar univariate general linear model (GLM). We introduce a MATLAB toolbox for fitting these models called Multivariate and Repeated Measures (MRM), comparing results from real neuroimaging datasets between this approach and other implementations of repeated-measures modelling of neuroimaging data. We also highlight the ability of this approach to integrating multimodal group-level imaging datasets. In addition, we discuss facilities in the MRM software to perform descriptive linear discriminant analysis (*d*LDA) to investigate how information from different modalities and sequences can be combined to maximally separate groups of interest. We also discuss the use of permutation-based approaches to *p*-value calculation, and multiple comparison corrections at both the voxel and cluster level, highlighting the utility of these methods when applied to the multivariate GLM.

2. Theory

The theory behind the multivariate extension of the univariate GLM is well documented (Christensen, 2001; Davis, 2002; Rencher and Christensen, 2012), and has recently been advocated for use in neuroimaging by Chen et al. (2014). Here we present a brief overview for completeness, emphasising how this approach is naturally adapted for repeated-measures/longitudinal models as well as multimodal integration. We also present the theory behind *d*LDA as an extension of the multivariate framework for understanding the contribution of multimodal imaging data to the separation of groups of interest.

2.1. The multivariate GLM

The multivariate form of the univariate GLM is expressed as

$$\mathbf{Y} = \mathbf{XB} + \mathbf{E} \tag{1}$$

where $\mathbf{Y}$ is an $n \times t$ matrix of observations, $\mathbf{X}$ is the $n \times k$ design matrix, $\mathbf{B}$ is the $k \times t$ matrix of model parameters, and $\mathbf{E}$ is the $n \times t$ matrix of errors. This can be written in matrix form as

$$\begin{pmatrix} Y_{11} & \cdots & Y_{1t} \\ \vdots & \cdot & \vdots \\ Y_{n1} & \cdots & Y_{nt} \end{pmatrix} = \begin{pmatrix} x_{11} & \cdots & x_{1k} \\ \vdots & \cdot & \vdots \\ x_{n1} & \cdots & x_{nk} \end{pmatrix} \begin{pmatrix} \beta_{11} & \cdots & \beta_{1t} \\ \vdots & \cdot & \vdots \\ \beta_{k1} & \cdots & \beta_{kt} \end{pmatrix} + \begin{pmatrix} \epsilon_{11} & \cdots & \epsilon_{1t} \\ \vdots & \cdot & \vdots \\ \epsilon_{n1} & \cdots & \epsilon_{nt} \end{pmatrix} \tag{2}$$

where n can be taken as the number of subjects, t as the number of *dependent variables*, here referred to as the repeated measurements or modalities, and k as the number of *independent variables*, here referred to as the predictors. Traditionally, it is assumed that $\boldsymbol{Y}_i \sim \mathcal{N}(\boldsymbol{X}_i\mathbf{B}, \boldsymbol{\Sigma})$ so that each ith row of $\mathbf{Y}$ is considered drawn from a multivariate normal distribution with a mean vector given by $\boldsymbol{X}_i\mathbf{B}$, and an unstructured covariance matrix $\boldsymbol{\Sigma}$. As with the univariate case, these assumptions can more usefully be expressed using the errors so that

$$\mathrm{Vec}(\boldsymbol{E}) \sim \mathcal{N}(0, \boldsymbol{I}_n \otimes \boldsymbol{\Sigma}) \tag{3}$$

where the Vec operator is used to re-express a matrix as a vector by stacking the transposed rows (Christensen, 2011; Rencher and Christensen, 2012). Here $\mathbf{0}$ is a vector of zeros, $\boldsymbol{I}_n$ is the $n \times n$ identity matrix, and $\otimes$ denotes the Kronecker product.

Estimation of $\mathbf{B}$ is usually performed using ordinary least squares,

$$\hat{\mathbf{B}} = (\boldsymbol{X}'\boldsymbol{X})^{-1}\mathbf{X}'\mathbf{Y} \tag{4}$$

identical to performing t univariate estimates using the columns of $\mathbf{Y}$. Here, the most salient difference with univariate approaches is evident as we no longer have a *vector* of estimated parameters but a *matrix*, with one column for each of the t dependent variables and one row for each of the k predictors in $\mathbf{X}$. Calculation of the multivariate residuals follows using $\hat{\mathbf{E}} = \mathbf{Y} - \mathbf{X}\hat{\mathbf{B}}$ so that an unbiased estimate of $\boldsymbol{\Sigma}$ can be made using

$$\hat{\boldsymbol{\Sigma}} = \frac{1}{n-k}\hat{\mathbf{E}}'\hat{\mathbf{E}} \tag{5}$$

(Davis, 2002; Rencher and Christensen, 2012). Here we see that the covariance structure of the model is both unconstrained and very simple to estimate. When applied to imaging data the residual matrix $\hat{\mathbf{E}}$ is estimated on a per-voxel basis and thus it is trivial to estimate a unique covariance structure for every voxel. This is a distinct advantage of mass multivariate approaches to dependent neuroimaging data. However, it should be clear from Eq. (3) that in this framework the covariance structure is assumed identical across groups. We shall return to this issue later.

The multivariate framework allows for the modelling of both repeated-measures and multimodal group-level imaging data. In both instances, each row of $\mathbf{Y}$ represents measurements from a single subject (for a particular voxel), with the columns of $\mathbf{Y}$ representing the multiple observations for that subject. Whether modelling repeated measurements or multiple modalities, there is an assumed degree of correlation between the columns of $\mathbf{Y}$. This correlation is expressed using the estimated variance–covariance matrix $\hat{\boldsymbol{\Sigma}}$, as indicated above. The utility of mixed-effects approaches for dependent data is in part due to their flexibility in specifying a variety of covariance structures (Mcculloch et al., 2008; Searle et al., 1992), whereas the assumption of a spherical covariance structure is one of the main reasons the traditional repeated-measures ANOVA approach is typically avoided (Davis, 2002). In the multivariate approach, an unconstrained covariance structure at every voxel provides the opportunity for inference without making any assumptions on the form that the covariance may take across the brain. As such, we argue that this is the safest approach without the computational burden of estimating variance components using iterative maximum-likelihood at every voxel (Guillaume et al., 2014). Notably, such a structure can also be fit uniquely at each voxel using marginal models, where the covariance structure is treated as a nuisance factor, allowing simplification of the mixed-effects scheme where both fixed and random effects must be specified directly (Guillaume et al., 2014; Li et al., 2013; Skup et al., 2012).

Extension of the multivariate GLM to accommodate continuous covariates is identical to the univariate domain and simply involves adding the, usually mean-centred (Poldrack et al., 2011), covariate w_i as another column in the design matrix $\mathbf{X}$. The parameters associated with w_i are therefore slopes of the relationship between w_i and $\mathbf{Y}$ for each column of $\mathbf{Y}$. If a grouping variable is used to split the covariate then a per-condition, or per-modality, slope is estimated for each group separately. Comparisons of changes in slope across groups are then easily specified. This scheme is more straightforward than integrating continuous covariates into traditional univariate approaches to repeated measurements, although it does not allow for the specification of time-varying covariates. With no groups and only continuous covariates the

[1] We use the term multimodal generically to cover both multiple sequences from the same imaging modality (e.g. fMRI, ASL, DTI) as well as the different imaging modalities themselves (e.g. MR, PET).

Please cite this article as: McFarquhar, M., et al., Multivariate and repeated measures (MRM): A new toolbox for dependent and multimodal group-level neuroimaging data, NeuroImage (2016), http://dx.doi.org/10.1016/j.neuroimage.2016.02.053