

ROI analysis of pharmafMRI data: an adaptive approach for global testing

Giorgos Minas, John A.D. Aston, Thomas E. Nichols and Nigel Stallard

Abstract Pharmacological fMRI (pharmafMRI) is a new highly innovative technique utilizing the power of functional Magnetic Resonance Imaging (fMRI) to study drug induced modulations of brain activity. FMRI recordings are very informative surrogate measures for brain activity but still very expensive and therefore pharmafMRI studies have typically small sample sizes. The high dimensionality of fMRI data and the arising high complexity requires sensitive statistical analysis in which often dimensionality reductions are crucial. We consider Region of Interest (ROI) analysis and propose an adaptive two-stage testing procedure for respectively formulating and testing the fundamental hypothesis as to whether the drug modulates the control brain activity in selected ROI. The proposed tests are proved to control the type I error rate and they are optimal in terms of the predicted chance of a true positive result at the end of the trial. Power analysis is performed by re-expressing the high dimensional domain of power function into a lower dimensional easily interpretable space which still gives a complete description of the power. Based on these results, we show under which circumstances our procedure outperforms standard single-stage and sequential two-stage procedures focusing on the small sample sizes typical in pharmafMRI. We also apply our methods to ROI data of a pharmafMRI study.

Key words: functional Magnetic Resonance Imaging, global testing, dimension reduction, adaptive designs, predictive power

Giorgos Minas

Department of Statistics and Warwick Centre of Analytical Sciences, University of Warwick, UK
e-mail: g.c.minas@warwick.ac.uk

John A.D. Aston

Department of Statistics, University of Warwick, UK e-mail: j.a.d.aston@warwick.ac.uk

Thomas E. Nichols

Department of Statistics and Warwick Manufacturing Group, University of Warwick, UK e-mail:
t.e.nichols@warwick.ac.uk

Nigel Stallard

Division of Health Sciences, Warwick Medical School, University of Warwick, UK e-mail:
n.stallard@warwick.ac.uk

1 Introduction

Pharmacological fMRI (pharmafMRI) is an exciting new technique employing functional Magnetic Resonance Imaging (fMRI) to study brain activity under drug administration. The so-called Blood Oxygenation Level Dependent (BOLD) fMRI contrast, often used in pharmafMRI studies, measures local blood flow changes known to be associated with changes in brain activity. While becoming more established, pharmafMRI faces a number of challenges of which some are statistical.

fMRI datasets are extremely high dimensional with enormous spatial resolution ($\approx 3mm$) and moderate temporal resolution ($\approx 3s$). The typical fMRI dataset produced by a single scanning session consists of BOLD recordings acquired during a relatively short period of time (few hundreds time points) from around 10^5 voxels (3-dimensional volume elements) throughout the brain. To handle such high dimensional datasets it is often appropriate to formulate specific regional hypotheses for the drug action and reduce accordingly the dimension of the data. The need for this type of analysis, which can provide regional summary measures of drug effect, is particularly acute in the typical pharmafMRI setting, in which due to the high cost of fMRI scans only a small number of subjects can be recruited.

Region of Interest (ROI) analysis can reduce an fMRI datasets into a relatively small number of ROI response summary measures expressing the local strength of treatment effect across the selected brain regions. If both the definition of ROI and the computation of the ROI response measures are cautiously conducted, a statistical analysis based on these ROI measures can potentially achieve high levels of sensitivity. We wish to go along this path and apply a multivariate test assessing the fundamental null hypothesis as to whether the new compound of interest changes the underlying brain activity in the selected ROI.

In previous work [5], we showed that tests based on a univariate linear combination of multivariate ROI responses can outperform fully multivariate methods, especially for the typically small sample sizes of fMRI studies. The decisive question for the former tests is the selection of the weights applied to ROI responses. O'Brien [6] used equal weights and Läuter [3] extract the weights from the data sums of products matrix while in Minas et al. [5] the weights are derived based on prior information and pilot data.

Here, we develop an adaptive two-stage procedure where a weighting vector, initially chosen based on prior information, is optimally adapted at a subsequent interim analysis based on the collected first stage data. The first and the second weighting vector are applied to the first and second stage responses, respectively, to produce the stage-wise linear combination test statistics. A combination function combining the test statistics of the two studies is used to perform the final analysis. Both weighting vectors are optimal in terms of the predictive power [7] of this two-stage test which is analytically proved to control the type I error rate.

Finally, we wish to assess the performance of our testing procedure. First note that the performance of a test with such a high dimensional design space can be hard to interpret. We tackle this problem by proving that the high dimensional domain of the power function can be re-expressed into a lower dimensional easily interpretable

space which still gives a complete description of the power. Using these results, our power analysis shows clearly those circumstances where our procedure outperforms standard single stage and two-stage sequential procedures. We also apply our methods to ROI data of a pharmafMRI study in which our tests are shown to be far more powerful than the latter methods.

2 Formulation

We start with a formal description of our problem.

Extraction of ROI measures: There are a number of different methods for extracting ROI summary measures from fMRI data. In all methods raw data preprocessing precedes modeling the series of fMRI 3-dim images typically using mass univariate General Linear Models (GLMs) at voxel-by-voxel resolution. Then, estimates of the treatment effect in each voxel of each subject are extracted from the GLM and averaged across the predefined, based on either brain anatomy or function, ROI. The coordinates of the produced multivariate outcome correspond to a representative measure of the treatment effect within each ROI of each subject.

Assumptions: We assume that the ROI responses of the n_j subjects participating in stage j of the study are independent multivariate Normal random variables

$$Y_{ji} \sim N_K(\mu, \Sigma), \quad i = 1, 2, \dots, n_j, \quad j = 1, 2, \quad (1)$$

with mean μ and covariance matrix Σ . Normality is typically an acceptable assumption for modeling ROI linear measures in fMRI [2].

Linear combinations: The ROI responses are summarized by scalar linear combinations

$$L_{ji} = \sum_{k=1}^K w_{jk} Y_{jik}, \quad (2)$$




Fig. 1 Typical steps of fMRI data analysis producing a multivariate ROI outcome. The preprocessed series of fMRI images are modeled at voxel-by-voxel resolution using mass univariate GLMs. Suitable estimates of parameter values (β) expressing the treatment effect in each voxel are first extracted from the GLM and then averaged across the predefined ROI.

where w_{jk} is the non-zero weight applied to the k -th ROI response, $k = 1, \dots, K$, of stage j .

Hypotheses: We wish to test the global null hypothesis of no treatment effect across all ROI $H_0 : \mu = \underline{0}$ ($= (0, 0, \dots, 0)^T$) against the two-sided alternative $H_1 : \mu \neq \underline{0}$.

Stage-wise statistics: The stage-wise linear combination z and t statistics for testing H_0 against H_1 when Σ is either known or unknown are respectively

$$Z_j = \frac{\bar{L}_j}{\sigma_j/n_j^{1/2}}, \quad T_j = \frac{\bar{L}_j}{s_j/n_j^{1/2}}. \quad (3)$$

Here, σ_j^2 , \bar{L}_j , s_j^2 are the variance, sample mean and sample variance of the linear combination L_j , respectively. The two-sided p values, p_j , $j = 1, 2$, may be obtained from the z or t statistics in (3).

Two-stage design: The two-stage design instructs the investigators to:

1. stop the trial (after the first stage) and reject H_0 if $p_1 < \alpha_1$ or stop the trial without rejection if $p_1 > \alpha_0$
2. continue to the second stage if $\alpha_1 \leq p_1 \leq \alpha_0$ and reject H_0 if $p_1 p_2 < c$.

Here, the Fisher's product combination function [1], $p_1 p_2$, is used for the final analysis. We also consider alternative functions including the Inverse Normal combination function [4].

Under this design, the **type I error rate** is controlled at the nominal α level if the rejection probability of the two-stage z or t test,

$$pr(p_1 < \alpha_1) + \int_{\alpha_1}^{\alpha_0} pr(p_1 p_2 < c | p_1) g(p_1) dp_1, \quad g(\cdot) \text{ density of } p_1, \quad (4)$$

is under the null hypothesis H_0 equal to α .

Target: We wish to maximise the **power** of the above two-stage tests, i.e. the rejection probability in (4) under H_1 , with respect to the weighting vectors w_1, w_2 , while controlling the type I error rate. In other words, we wish to find the optimal direction in which the projection of the treatment effect vector produces optimal power.

3 Methods

Here, we develop the proposed adaptive two-stage testing procedure. We start by providing the optimal weighting vector for the two-stage z and t tests described

above.

Theorem 1. *Under the assumption in (1), the power of the above two stage tests, i.e. the rejection probability in (4) under H_1 , is **maximized** with respect to w_1 and w_2 if and only if the latter are both proportional to $\omega = \Sigma^{-1}\mu$.*

The optimal vector ω is unknown and therefore we use the available information at the planning stage (prior) and at the interim stage (posterior) to select w_1, w_2 .

Prior: Prior information D_0 elicited from previous studies and experts clinical opinion is used to inform the following Normal and Inverse-Wishart priors for μ and Σ , respectively,

$$(\mu \mid \Sigma, D_0) \sim N_K(m_0, \Sigma/n_0), \quad (\Sigma \mid D_0) \sim IW_{K \times K}(v_0, S_0^{-1}). \quad (5)$$

Here, m_0 represents a prior estimate for μ , n_0 the number of observations m_0 is based on; and v_0, S_0 respectively represent the degrees of freedom and scale matrix of the inverse-Wishart prior.

Posterior: Under this Bayesian model, the posterior distributions, given the prior information D_0 and the first stage data y_1 , have the same form as the prior distributions

$$(\mu \mid \Sigma, D_0, y_1) \sim N_K(m_1, \Sigma/(n_0 + n_1)), \quad (\Sigma \mid D_0, y_1) \sim IW_{K \times K}(v_0 + n_1, S_1^{-1}). \quad (6)$$

where the posterior estimates

$$m_1 = \frac{n_0 m_0 + n_1 \bar{y}_1}{n_0 + n_1}, \quad S_1 = S_0 + (n_1 - 1)S_{y_1} + \frac{n_0 n_1}{n_0 + n_1} (\bar{y}_1 - m_0)(\bar{y}_1 - m_0)^T \quad (7)$$

can be thought as “weighted averages” of the prior and first stage estimates of μ and Σ , respectively.

The predictive power of our test, which we are interested in maximizing with respect to the weighting vectors w_1, w_2 , is introduced next.

Predictive power: Predictive power expresses “*the chance, given the data so far, that the planned test rejects H_0 when the trial is completed*”. Given D_0 , the predictive power $B_{z,1}$ and $B_{t,1}$ for the two-stage z and t tests, respectively, are equal to

$$pr(p_1 < \alpha_1 \mid D_0) + pr(p_1 \in [\alpha_1, \alpha_0], p_1 p_2 < c \mid D_0) \quad (8)$$

and if we continue to the second stage, the predictive power $B_{z,2}$ and $B_{t,2}$ given the prior information D_0 and the first stage data \mathbf{y}_1 are equal to

$$pr(p_1 p_2 < c \mid D_0, \mathbf{y}_1), \quad (9)$$

for p_j corresponding to either the z or t statistics in (3), respectively.

Theorem 2. *Under the assumptions (1) and (6), the first and second stage predictive power of the z test, $B_{z,1}$ and $B_{z,2}$ are maximized with respect to w_1 , w_2 , respectively, if the latter are proportional to $w_{z,1} = \Sigma^{-1}m_0$ and $w_{z,2} = \Sigma^{-1}m_1$, respectively.*

Further, for large v_0 , i.e. $v_0 \rightarrow \infty$, the weighting vectors $w_{t,1} = S_0^{-1}m_0$ and $w_{t,2} = S_1^{-1}m_1$ maximise the predictive power functions $B_{t,0}$ and $B_{t,1}$.

Adaptive two-stage z and t tests: The proposed tests follow the two-stage design described earlier with the first and second stage weighting vectors of the stage-wise z and t statistics being equal to $w_{z,1}, w_{z,2}$ and $w_{t,1}, w_{t,2}$, respectively. These tests are power optimal based on the collected information. We can also prove that they control the type I error rate.

4 Power analysis

The design variables that need to be considered for the analysis of the power function of the above z and t tests are: (i) the stopping boundaries α_0 , α_1 and c , (ii) the sample sizes n_0 , n_1 and n_2 (and v_0), (iii) the parameters μ and Σ and (iv) the prior estimate(s) m_0 (and S_0). While the variables in (i) and (ii) are scalar, those in (iii) and (iv) are high dimensional ($\mathbb{R}^K \times \mathbb{R}^K \times \mathbb{R}^{K \times K} (\times \mathbb{R}^{K \times K})$). Without any reduction, it would be challenging to get a full picture and explain the power performance of our tests. However, we prove that for the z test, (iii) and (iv) can be replaced by: (a) the Mahalanobis distance $(\mu \Sigma^{-1} \mu)^{1/2}$ of the null $N_K(\mathbf{0}, \Sigma)$ to the alternative $N_K(\mu, \Sigma)$ distribution expressing the strength of the treatment effect and (b) the angle θ between the selected and the optimal direction. The same stands for the t test although the angular distance in (b) is replaced by one expressed in terms of easily interpretable vectors in $[0, \pi/2]^K \times [0, \pi/2]^K \times \mathbb{R}^K$. In figure 1, we illustrate how these results can be used to compare our procedure to standard testing procedures. For each set values of the parameters in (i)-(ii), there is a fixed value of the angle θ for which smaller values of θ result in higher power for the z and t tests compared to the classical multivariate tests and vice versa.

4.1 Application to a pharmafMRI study:

We use the sample mean and sample covariance matrix (see table 1) of ROI data extracted from a GlaxoSmithKline pharmafMRI study ($K = 11$, $n_T = 13$) to perform power comparisons. Even for such small sample sizes and the fairly poor prior estimates in table 1, the adaptive t test might be considered as sufficiently powered ($\beta_t = 0.82$). This is in contrast to standard single stage tests, such as Hotelling's T^2 and OLS tests [6] ($\beta_{T^2} = 0.30$, $\beta_{OLS} = 0.13$) as well as their corresponding sequential two-stage versions (sequential Hotelling's T^2 test not applicable for $n_T = 13$, $\beta_{OLS}^s = 0.10$).




Fig. 2 Simulation-based approximation of the power, β_t , of the single-stage (green —) and adaptive (blue —) linear combination t test as well as the Hotelling's T^2 test (red $\cdot\cdot$) plotted against the total sample size n_T . The angle θ between the optimal weighting vector and the selected weighting vector either of the single-stage t test or the adaptive t test is taken to be equal to 0° (*), 15° (\circ), 30° (\square), 45° (∇), 60° (\times), 75° ($+$) and 90° (\star). For small sample sizes, the power of the single-stage t test is larger (smaller) than the power of the adaptive t test if the angle between the optimal and the selected weighting vector is small (large). For relatively large sample sizes, in contrast to the single stage test, the adaptive t test reaches high power levels even for first stage weighting vector orthogonal (90°) to the optimal. For increasing n_T and all other design variables remaining fixed, the angle θ^* , for which the power of the T^2 test (applicable only for $n_T > K$) is equal to the power of the t test, is decreasing. Here $\alpha_0 = 1$, $\alpha_1 = 0.01$, $c = 0.0087$ ($\alpha = 0.05$), $n_0 = 5$, $v_0 = 4$, $f = n_1/n_T = 0.5$, $D_1 = \Sigma^{-1/2}S_0\Sigma^{-1/2} = I$, $K = 15$.

5 Discussion

The formulation of specific regional hypotheses for drug action and the associated dimensionality reductions are crucial for further establishment of pharmafMRI. As we illustrate in our methods, ROI analysis combined with multivariate methods can be used to answer the fundamental question as to whether the drug modulates the brain activity over the regions of greatest interest for the study. We show that reducing the ROI responses into a univariate linear combination can be advantageous in terms of sensitivity compared to fully multivariate methods on ROI responses, without any cost in terms of specificity. For the latter dimensionality reduction, we propose deriving the weights of the linear combination by exploiting the available prior information and allowing for data dependent adaptation at an interim analysis. These weights are optimal in terms of predictive power give the available information at each selection time. Further, we show how the high dimensional power function domain space is reduced substantially to a lower dimensional easily interpretable space which allows us to show clearly under which circumstances the improvement over single stage and sequential designs is achieved. We finally show

Table 1 Means (line 1), variances (line 3) and correlations (upper triangle of matrix in lines 5 – 15) and the corresponding prior estimates (lines 2, 4 and lower triangle of matrix in lines 5 – 15) of ROI data of the sample ($n_T = 13$) of a GSK pharmafMRI study. The ROI are: *Anterior Cingulate (AC)*, *Atlas Amygdala (A)*, *Caudate (C)*, *Dorsolateral Prefrontal Cortex (DLPFC)*, *Globus Pallidus (GP)*, *Insula (I)*, *Orbitofrontal cortex (OFC)*, *Putamen (P)*, *Substantia Nigra (SA)*, *Thalamus (T)*, *Ventral Striatum (VS)*. Effect sizes differ across ROI and generally high correlations are observed. Despite the fairly large distance (angle 67°) between the optimal weighting vector and the first stage weighting vector (for the prior estimates in the table below) the adaptive t test for $\alpha_0 = 1$, $\alpha_1 = 0.01$, $c = 0.0087$ ($\alpha = 0.05$), $n_0 = 3$, $n_1 = 6$, $n_2 = 7$ might be considered as sufficiently powered ($\beta_t = 0.82$).

	ROI	AC	A	C	DLPFC	GP	I	OFC	P	SA	T	VS
1	μ_k	-0.01	0.06	-0.08	-0.08	-0.14	-0.02	-0.08	-0.06	-0.10	-0.10	-0.13
2	$m_{0,k}$	0	0.10	-0.10	-0.10	-0.15	0	-0.15	0	-0.10	-0.10	-0.15
3	σ_k	0.11	0.11	0.03	0.05	0.11	0.08	0.13	0.15	0.10	0.11	0.10
4	$s_{0,k}$	0.15	0.10	0.02	0.10	0.10	0.10	0.15	0.15	0.10	0.10	0.10
5	AC	1	0.70	0.87	0.88	0.73	0.89	0.66	0.81	0.26	0.95	0.70
6	A	0.70	1	0.54	0.61	0.72	0.77	0.65	0.68	0.59	0.68	0.66
7	C	0.70	0.50	1	0.89	0.72	0.87	0.47	0.80	0.27	0.90	0.74
8	DLPFC	0.70	0.70	0.70	1	0.71	0.76	0.73	0.77	0.27	0.87	0.62
9	GP	0.70	0.70	0.70	0.70	1	0.86	0.51	0.90	0.54	0.70	0.90
10	I	0.70	0.70	0.70	0.70	0.70	1	0.45	0.85	0.46	0.86	0.84
11	OFC	0.50	0.50	0.50	0.70	0.50	0.50	1	0.44	0.09	0.65	0.30
12	P	0.70	0.70	0.70	0.70	0.70	0.70	0.50	1	0.49	0.82	0.89
13	SA	0.50	0.70	0.30	0.50	0.50	0.50	0.50	0.30	1	0.30	0.55
14	T	0.70	0.70	0.70	0.70	0.70	0.70	0.50	0.70	0.50	1	0.74
15	VS	0.70	0.50	0.70	0.70	0.70	0.70	0.50	0.70	0.50	0.70	1

that our methods can outperform standard single stage and sequential two-stage multivariate tests in a pharmafMRI study.

References

1. P. Bauer and K. Kohne. Evaluation of Experiments with Adaptive Interim Analyses. *Biometrics*, 50(4):pp. 1029–1041, 1994.
2. K. Friston, J. Ashburner, S. Kiebel, T. Nichols, and W. Penny. *Statistical parametric mapping : the analysis of functional brain images*. Elsevier/Academic Press, Amsterdam; Boston, 2007.
3. J. Lauter. Exact t and F tests for analyzing studies with multiple endpoints. *Biometrics*, 52:964–970, 1996.
4. W. Lehmacher and G. Wassmer. Adaptive Sample Size Calculations in Group Sequential Trials. *Biometrics*, 55(4):1286–1290, 1999.
5. G. Minas, F. Rigat, T.E. Nichols, J.A.D. Aston, and N. Stallard. A hybrid procedure for detecting global treatment effects in multivariate clinical trials: theory and applications to fMRI studies. *Statist. Med.*, 31(3):253–268, 2012.
6. P. C. O’Brien. Procedures for comparing samples with multiple endpoints. *Biometrics*, 40:1079–1087, 1984.
7. D. J. Spiegelhalter, L. S. Freedman, and P. Blackburn. Monitoring clinical trials: conditional or predictive power? *Controlled Clinical Trials*, 7:8–17, 1986.