

Faster family-wise error control for neuroimaging with a parametric bootstrap

SIMON N. VANDEKAR*

Department of Biostatistics, Epidemiology, and Informatics, 423 Guardian Dr., University of Pennsylvania, Philadelphia PA 19104, USA

simonv@mail.med.upenn.edu

THEODORE D. SATTERTHWAITE, ADON ROSEN, RASTKO CIRIC, DAVID R. ROALF,
KOSHA RUPAREL, RUBEN C. GUR, RAQUEL E. GUR

Department of Psychiatry, University of Pennsylvania, Philadelphia, PA 19104, USA

RUSSELL T. SHINOHARA

Department of Biostatistics, Epidemiology, and Informatics, 423 Guardian Dr., University of Pennsylvania, Philadelphia PA 19104, USA

SUMMARY

In neuroimaging, hundreds to hundreds of thousands of tests are performed across a set of brain regions or all locations in an image. Recent studies have shown that the most common family-wise error (FWE) controlling procedures in imaging, which rely on classical mathematical inequalities or Gaussian random field theory, yield FWE rates (FWER) that are far from the nominal level. Depending on the approach used, the FWER can be exceedingly small or grossly inflated. Given the widespread use of neuroimaging as a tool for understanding neurological and psychiatric disorders, it is imperative that reliable multiple testing procedures are available. To our knowledge, only permutation joint testing procedures have been shown to reliably control the FWER at the nominal level. However, these procedures are computationally intensive due to the increasingly available large sample sizes and dimensionality of the images, and analyses can take days to complete. Here, we develop a parametric bootstrap joint testing procedure. The parametric bootstrap procedure works directly with the test statistics, which leads to much faster estimation of adjusted p -values than resampling-based procedures while reliably controlling the FWER in sample sizes available in many neuroimaging studies. We demonstrate that the procedure controls the FWER in finite samples using simulations, and present region- and voxel-wise analyses to test for sex differences in developmental trajectories of cerebral blood flow.

Keywords: Hypothesis testing; FWE control; Neuroimaging.

*To whom correspondence should be addressed.

1. INTRODUCTION

Magnetic resonance imaging (MRI) is a widely used tool for studying the neurological correlates of human cognition, psychiatric disorders, and neurological diseases. This is due to the flexibility of MRI to noninvasively study various functional and physiological properties of the human brain. Often, many hypothesis tests are performed at every voxel or at anatomically defined brain regions in order to identify locations that are associated with a cognitive or diagnostic variable. Multiple testing procedures (MTPs) are crucial for controlling the number of false positive findings within a statistical parametric map or across a set of brain regions being investigated. Typically the family-wise error rate (FWER) is controlled at a level $0 < \alpha < 1$, meaning that the probability one or more null hypotheses is falsely rejected is less than or equal to α .

Recently, several studies have demonstrated that commonly used FWER controlling procedures yield incorrect false-positive rates (Eklund *and others*, 2016, 2012; Silver *and others*, 2011). Cluster-based spatial inference procedures (Friston *and others*, 1994) that rely on Gaussian random field (GRF) theory can have hugely inflated false-positive rates, while voxel-wise GRF MTPs (Friston *and others*, 1994) tend to have exceedingly small FWERs that are far below the nominal level. The failure of GRF procedures is due to the fact that the spatial assumptions of GRF approaches are often violated in neuroimaging data sets (Eklund *and others*, 2016). The small type 1 error rate of voxel-wise procedures is due to the reliance on classical FWER procedures (such as the Bonferroni procedure) that do not account for the strong dependence between hypothesis tests in voxel-wise and region-wise analyses. This small type 1 error rate leads to an inflated type 2 error rate and loss of power. These recent studies demonstrate a dire need for robust and powerful inference procedures.

To our knowledge, the only methods used in neuroimaging that reliably control the FWER are permutation-based joint testing procedures (Winkler *and others*, 2014; Eklund *and others*, 2016; Dudoit *and van der Laan*, 2008). These methods maintain the nominal FWER because they appropriately reproduce the joint distribution of the imaging data, thereby overcoming the limitations of methods typically used in imaging. Unfortunately, permutation testing is computationally intensive, especially in modern imaging data sets that have large sample sizes and hundreds of thousands of voxels or hundreds of brain regions. Moreover, currently available neuroimaging software only performs single-step testing procedures, although step-down procedures are uniformly more powerful. The extensive computation time means it can take days to perform statistical tests or can even lead investigators to reduce the number of permutations to an extent that adjusted p -values have large error.

As a solution, we design a parametric-bootstrap joint (PBJ) testing procedure for hypothesis testing in region- and voxel-wise analyses. Region-wise analyses are performed by averaging all voxel values within anatomically defined regions and fitting a model at each region. Voxel-wise analyses fit a model at each of hundreds of thousands of voxels in a brain image. As the parametric bootstrap does not require resampling and refitting the model for every iteration, it is faster than permutation testing procedures. In addition, our procedure allows the generated null distribution to be applied to multiple tests of statistical parameters. This drastically reduces computing time as the null distribution can be estimated from one bootstrap procedure and applied for many tests. We demonstrate the efficacy of our procedure by investigating sex differences in development-related changes of cerebral blood flow (CBF) measured using arterial spin labeled MRI (Satterthwaite *and others*, 2014).

In Section 2, we discuss several FWER controlling procedures used in neuroimaging and classify them with regard to single-step/step-down and marginal/joint procedures. In Section 3, we present the new PBJ procedure. We summarize the data and analyses in Section 4 and the simulation methods in Section S1.2 of the supplementary material available at Biostatistics online. In Section 5, we use simulations to investigate when joint MTPs maintain the nominal type 1 error rate, and we compare the power and FWER of the PBJ to commonly used MTPs using simulations of region- and voxel-wise data analyses. Finally, in Section 6 we perform region- and voxel-wise analyses of the CBF data.

2. OVERVIEW OF MULTIPLE TESTING PROCEDURES

Throughout, we will assume the image intensity for n subjects, $Y_v \in \mathbb{R}^n$, for voxels or regions, $v = 1, \dots, V$, can be expressed as the linear model

$$Y_v = X_0\alpha_v + X_1\beta_v + \epsilon_v = X\zeta_v + \epsilon_v, \quad (2.1)$$

where X_0 is an $n \times m_0$ matrix of nuisance covariates, X_1 is an $n \times m_1$ matrix of variables to be tested, $m = m_0 + m_1$, $X = [X_0, X_1]$, parameters $\alpha_v \in \mathbb{R}^{m_0}$, $\beta_v \in \mathbb{R}^{m_1}$, and $\zeta_v = [\alpha_v^T, \beta_v^T]^T$. Let $Y = [Y_1, \dots, Y_V]$ and let Y_i denote an arbitrary row vector of Y . Assume that the $V \times V$ covariance matrix is the same for each subject, $\text{cov}(Y_i) = \Psi$, and define the correlation matrix

$$\Sigma_{j,k} = \Psi_{j,k} / \sqrt{\Psi_{j,j}\Psi_{k,k}}. \quad (2.2)$$

We denote the observed test statistics by Z_{v0} for $v = 1, \dots, V$, where we reject the null $H_0 : \beta_v = 0$ for large values of Z_{v0} . The notation Z is used (as opposed to F) as we consider transformed F-statistics in Section 3.

At each location we are interested in performing the test of the null hypothesis

$$H_{0v} : \beta_v = 0$$

using an F-statistic. The form of model (2.1) covers a wide-range of possible tests including tests of group differences, continuous covariates, analysis of variance, and interactions. V is typically in the hundreds for region-wise analyses or the hundreds of thousands for voxel-wise analyses.

The goal of all MTPs is to control some measure of the number of false positive findings in a family of hypothesis tests. We will assume the approach of controlling the FWER at some level $0 < \alpha < 1$. In most fields, we would like to maintain control of the FWER even in the case that there are false null hypotheses. This is referred to as strong control of the FWER (Hochberg, 1988).

DEFINITION 2.1 Let $\{H_1, \dots, H_V\} = H$ denote a set of hypotheses. A correction procedure has α level strong control of the FWER if for all $H' \subset H$ of true null hypotheses

$$\mathbb{P}(\text{retain } H_v \text{ for all } H_v \in H') \geq 1 - \alpha. \quad (2.3)$$

In neuroimaging, strong control of the FWER corresponds to maintaining the correct FWER control even if there is a set of regions or voxels where there is a true effect. The strong FWER controlling procedures discussed in this manuscript are given in Table S1 of the supplementary material available at Biostatistics online.

2.1. Single-step and Step-down procedures

Due to the complex dependence structure of the test statistics in neuroimaging data, it is necessary to use testing procedures that are appropriate for any type of dependence among tests. Single-step and step-down procedures are two classes of MTPs that have strong control of the FWER for any dependence structure (Dudoit and others, 2003). These are in contrast to step-up procedures, which make explicit assumptions about the dependence structure of the test statistics (Sarkar and Chang, 1997). When testing V hypotheses $H = \{H_1, \dots, H_V\}$ in a family of tests, the single-step procedures use a more stringent common threshold $\alpha^* \leq \alpha$ such that the inequality (2.3) is guaranteed. While this procedure is simple, in most cases it is uniformly more powerful to use a step-down procedure.

PROCEDURE 2.1 (Step-down procedure) Let $p_{(1)}, \dots, p_{(V)}$ be the increasingly ordered p -values, for hypotheses $H_{(1)}, \dots, H_{(V)}$. The step-down procedure uses thresholds $\alpha_{(1)}^* \leq \dots \leq \alpha_{(V)}^* \leq \alpha$ to find the largest value of K such that

$$p_{(k)} < \alpha_{(k)}^* \text{ for all } k \leq K,$$

and rejects all hypotheses $H_{(1)}, \dots, H_{(K)}$.

The single-step procedure can usually be improved by modifying it to be step-down procedure while still maintaining strong control of the FWER (Holm, 1979; Marcus and others, 1976; Hommel, 1988). The canonical example of this modification involves the Bonferroni and Holm procedures (Dunn, 1961; Holm, 1979). The Bonferroni procedure uses the common threshold $\alpha^* = \alpha/V$ for all hypotheses and rejects all H_k such that $p_k < \alpha/V$. The Holm procedure instead uses the thresholds $\alpha_{(k)}^* = \alpha/(V + 1 - k)$ and rejects using Procedure 2.1. Holm's procedure is always more powerful than Bonferroni and still controls the FWER strongly (Holm, 1979).

2.2. Joint testing procedures

MTPs can further be classified into marginal and joint testing procedures. The Bonferroni and Holm approaches are called marginal procedures because they do not make use of the dependence of the test statistics. As they must be able to control any dependence structure, they are more conservative than joint testing procedures that cater exactly to the distribution of the test statistics (Dudoit and van der Laan, 2008). The benefit of accounting for the dependence of test statistics is critical in neuroimaging, where the test statistics are highly dependent due to spatial, anatomical, and functional dependence. Joint MTPs differ in how the null distribution is estimated from the sample, but use the same procedure to compute single-step or step-down adjusted p -values. For this reason, we will first discuss the estimation of the null distribution and then discuss how adjusted p -values are computed from the estimate of the null.

2.2.1. *Estimating the null distribution with permutations* The ‘‘Randomise’’ procedure proposed by Winkler and others (2014) is a single-step permutation joint (PJ) MTP widely used in neuroimaging. The PJ MTP procedure is a modification of the Freedman–Lane procedure (Freedman and Lane, 1983) implemented by Winkler and others (2014) that estimates the null distribution of the test statistics by permuting the residuals of the reduced model to obtain estimates of the parameters under the null hypothesis that there is no association between the variables in X_1 and the outcome. Though only the single-step procedure has been proposed for use in neuroimaging, for completeness we include null estimation of the test statistics for the step-down procedure.

PROCEDURE 2.2 (Permutation null estimation) Assuming the model (2.1):

1. Regress Y_v against the reduced model $Y_v = X_0\alpha_v + \epsilon_v$ to obtain the residuals $\hat{\epsilon}_v$ and the test statistics Z_{v0} for all regions or voxels $v = 1, \dots, V$.
2. Order the test statistics $Z_{(1)0} < Z_{(2)0} < \dots < Z_{(V)0}$ and let $\epsilon_{(v)}$ be the corresponding residuals.
3. For $b = 1, \dots, B$, randomly generate a permutation matrix P_b , permute the residuals $\hat{\epsilon}_{(v)b} = P_b\hat{\epsilon}_{(v)}$, and define the permuted data at each voxel as $Y_{(v)b} = \hat{\epsilon}_{(v)b}$.
4. For $v = 1, \dots, V$ and $b = 1, \dots, B$ regress $Y_{(v)b}$ onto the full model (2.1) to obtain the test statistic $Z_{(v)b}$ to be used as an estimate of the null distribution.

Ordering the test statistics is not necessary for the single-step procedure, but is required for computing step-down adjusted p -values. Note that for any given b , the generated test statistics $Z_{(v)b}$ may not be increasing in v . Strong control of the FWER for this permutation procedure relies on the assumption of subset pivotality (see Supplementary Material available at *Biostatistics* online) (Westfall and Young, 1993).

This null distribution can be used to compute rejection regions or adjusted p -values. Because all joint testing procedures rely on estimates of the null distribution, they are approximate in finite samples. The permutation methodology and our proposed PBJ procedure (Section 3) only differ in how the null distribution is estimated.

2.2.2. Computing adjusted p -values The following procedures describe how to obtain single-step and step-down adjusted p -values using any estimate of the null distribution generated by permutation or bootstrapping.

PROCEDURE 2.3 (Single-step joint adjusted p -values) Assuming the model (2.1) and an empirical distribution of null statistics $Z_{(v)b}$ for $v = 1, \dots, V$ and $b = 1, \dots, B$:

1. Compute $Z_{\max, b} = \max_{v \leq V} Z_{(v)b}$.
2. Compute the voxel-wise corrected p -value as $\tilde{p}_v = 1/B \sum_{b=1}^B I(Z_{\max, b} \geq Z_{(v)0})$, where $I(\cdot)$ is the indicator function.

PROCEDURE 2.4 (Step-down joint adjusted p -values) Assuming the model (2.1) and an empirical distribution of null statistics $Z_{(v)b}$ for $v = 1, \dots, V$ and $b = 1, \dots, B$:

1. Compute the statistics $Z_{\max v, b} = \max_{k \leq v} Z_{(k)b}$.
2. Compute the intermediate value $p_{(v)}^* = 1/B \sum_{b=1}^B I(Z_{\max v, b} \geq Z_{(v)0})$.
3. The adjusted p -values are $\tilde{p}_{(v)} = \max_{k \leq v} p_{(k)}^*$.

The single-step procedure is less powerful than the step-down counterpart (Dudoit and van der Laan, 2008). This is evident in comparing the procedures, as the adjusted p -values for the step-down procedure are at least as small as the adjusted p -values from the single-step. The adjusted p -values obtained using the step-down approach correspond to using Procedure 2.1. The key feature of Procedure 2.3 is that the estimate of the joint distribution is used to compute adjusted p -values, whereas Holm's procedure is a version of Procedure 2.1 that only uses the marginal distribution of the test statistics. So far, we have described existing MTPs used in neuroimaging.

3. PARAMETRIC-BOOTSTRAP

In this section, we propose single-step and step-down PBJ approaches that are conceptually identical to the PJ procedure, but differ in how the null distribution of the statistics is generated. The PBJ is based on the theory developed by Dudoit and van der Laan (2008) and therefore does not rely on the assumption of subset pivotality. We will allow the additional assumption that under the null the test statistics are approximately chi-squared. The chi-squared approximation can rely on asymptotic results, or as we show in Section 4.2, a transformation can be used so that the test statistics are approximately chi-squared. As with the PJ procedure discussed above this implies that the p -values are approximations that become more accurate as $n \rightarrow \infty$. In Section 5, we use simulations to show that the procedures control the FWER in sample sizes available in many neuroimaging studies.

3.1. Asymptotic control of the FWER

Here, we give a brief overview of the underlying assumptions sufficient to prove that the adjusted p -values from the PBJ control the FWER asymptotically. Details are given in the Supplementary Material available at *Biostatistics* online. We require that the test statistics' null distribution satisfies the null domination condition (Dudoit and van der Laan, 2008, p. 203) and need a consistent estimate of the null distribution.

DEFINITION 3.1 (Asymptotic null domination) Let H_0 denote the indices of M true null hypotheses in the set of V hypotheses $H = \{H_1, \dots, H_V\}$, with corresponding test statistics Z_{1n}, \dots, Z_{Vn} . The V -dimensional null distribution Q_0 satisfies the *asymptotic null domination* condition if for all $x \in \mathbb{R}$

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left(\max_{m \in H_0} Z_{mn} > x \right) \leq \mathbb{P} \left(\max_{m \in H_0} Z_m > x \right),$$

where $Z_n \sim Q_n$ is distributed according to a finite sample null joint distribution Q_n and $Z \sim Q_0$.

The joint null distribution Q_0 for the test statistics can be used to compute asymptotically accurate adjusted p -values if the null domination condition holds. We use a diagonal singular Wishart distribution as the null because it is proportional to the asymptotic distribution of a vector of F-statistics. We also transform the F-statistics marginally to chi-squared statistics if

$$\epsilon_v \sim \mathcal{N}(0, \sigma_v I), \quad (3.4)$$

for the error term in model (2.1). The assumption of asymptotic null domination of Definition 3.1 is satisfied when using our transformed F-statistics even if the error distribution is not normal (see Theorem S2 in the Supplementary material available at *Biostatistics* online). Thus, the PBJ provides approximate p -values regardless of the error distribution.

In practice the joint distribution of the test statistics, Q_0 , is not known *a priori* and must be estimated from the data. In order to obtain asymptotically valid adjusted p -values, the estimate for Q_0 , \hat{Q}_0 , must be consistent. Because the distribution $Q_0 = Q_0(\Sigma)$ is a function of the covariance matrix Σ , a consistent estimator for Q_0 can be obtained from a consistent estimator for Σ . The choice of the estimator is critical because $\hat{\Sigma}$ must be consistent for Σ under the alternative distribution. That is, even if some null hypotheses are false $\hat{\Sigma}$ must be consistent. The PBJ procedure uses a consistent estimator for Σ based on the residuals of the full design X . The consistency of $\hat{\Sigma}$ under the alternative guarantees asymptotic control of the FWER (see Supplementary Material available at *Biostatistics* online). Note that, in general, the PJ procedure may not yield a consistent estimator for the joint distribution Q_0 if the alternative is true at more than one location because the estimates of covariances are biased. The covariance estimates are biased due to the fact that, for the reduced model used by the PJ MTP (see Procedure 2.2), the mean is incorrectly specified in locations where the alternative is true. If the assumption of subset pivotality is satisfied, then the permutation estimator will be consistent.

3.2. Parametric bootstrap null distribution

For the parametric bootstrap we assume model (2.1) and use F-statistics for the test $H_{0v} : \beta_v = 0$.

$$F_{vn} = \frac{(n-m)Y_v^T(R_{X_0} - R_X)Y_v}{m_1 Y_v^T R_X Y_v}, \quad (3.5)$$

where R_A denotes the residual forming matrix for A . When the errors are normally distributed (3.4), F_{vn} is an F-distributed random variable with m_1 and $n-m$ (DOF). When the errors are not normal the statistics

(3.5) are asymptotically $m_1^{-1}\chi_{m_1}^2$. The following theorem gives the asymptotic joint distribution of the statistics.

THEOREM 1 Assume model (2.1), let F_{vn} be as defined in (3.5), and define the $p \times V$ matrix $\alpha = [\alpha_1, \dots, \alpha_V]$. Further assume that, under the null,

$$R_{X_0} \mathbb{E}Y = \mathbb{E}Y - X_0\alpha = 0 \quad (3.6)$$

$$\|\Psi\|_M < \infty, \quad (3.7)$$

where $\|\Psi\|_M = \sup_x \|\Psi x\|/\|x\|$ is the natural norm.

Then the following hold:

1. Define the matrix $\Phi_{i,i} = 1/\sqrt{\Psi_{i,i}}$ and $\Phi_{i,j} = 0$ for $i \neq j$. When (3.4) holds,

$$\begin{aligned} \Phi Y^T (R_{X_0} - R_X) Y \Phi &\sim \mathcal{W}_V(m_1, \Sigma) \\ \Phi Y^T R_X Y \Phi &\sim \mathcal{W}_V(n - m, \Sigma), \end{aligned} \quad (3.8)$$

where $\mathcal{W}_p(d, \Sigma)$ denotes a singular Wishart distribution with (DOF) $d < p$ and matrix $\Sigma \in \mathbb{R}^{p \times p}$ (Srivastava, 2003).

2. The F-statistics converge in law to the diagonal of a singular Wishart distribution, that is,

$$m_1 F_n = m_1 [F_{1n}, \dots, F_{Vn}] \rightarrow_L \text{diag} \{ \mathcal{W}_V(m_1, \Sigma) \}.$$

as $n \rightarrow \infty$.

In order to make the statistics robust regardless of whether the errors are normal, we use the transformation

$$Z_{vn} = \Phi^{-1} \{ \Phi_n(F_{vn}) \}, \quad (3.9)$$

where Φ_n is the cumulative distribution function (CDF) for a $\mathcal{F}(m_1, n - m)$ random variable and Φ^{-1} is the inverse CDF for a $\chi_{m_1}^2$ random variable. If the errors are normal (3.4), then the transformed statistics (3.9) are marginally $\chi_{m_1}^2$ statistics. When the errors are nonnormal then Z_{vn} is asymptotically $\chi_{m_1}^2$ (see Theorem S2 in the supplementary material available at Biostatistics online). The asymptotic joint distribution of the statistics given in Theorem 1 allows us to use a diagonal singular Wishart to compute approximate adjusted p -values. To compute probabilities we do not have to sample the full matrix (3.8), since only the diagonal elements, $\text{diag}\{Y^T(R_X - R_{X_F})Y\}$, are required.

To find adjusted p -values, \tilde{p}_v , for the single-step procedure we compute the probability

$$\tilde{p}_v = \mathbb{P} \left(\max_{k \leq V} |Z_k| > |Z_{v0}| \right), \quad (3.10)$$

where

$$Z = (Z_1, \dots, Z_V) \sim \text{diag} \{ \mathcal{W}_V(m_1, \Sigma) \}, \quad (3.11)$$

and Z_{v0} is the observed statistic at location v . Theorems S2 and S4 in supplementary material available at Biostatistics online guarantee asymptotic control of the FWER when (3.11) is used as the null distribution.

In practice, the joint distribution (3.11) is unknown due to the fact that Σ is unobserved, so the probability (3.10) cannot be computed. We must obtain an estimate for Σ in order to compute estimates of these probabilities. Since the diagonal, $\Sigma_{v,v} = 1$, we only need to estimate the off-diagonal elements. By estimating ρ_{ij} with the consistent estimator

$$\hat{\rho}_{jk} = \frac{Y_j^T R_X Y_k}{\hat{\sigma}_j \hat{\sigma}_k},$$

we are guaranteed asymptotic control of the FWER (see Supplementary Material available at Biostatistics online). This estimator is biased toward zero in finite samples, and yields conservative estimates of the correlation. Note that using the residuals of the full model is crucial here as that estimator yields consistent estimates of the correlation regardless of whether the alternative is true in each location. Instead of using Σ in (3.11) we use the estimated covariance matrix of the test statistics

$$\hat{\Sigma}_{jk} = \begin{cases} 1 & : j = k \\ \hat{\rho}_{jk} & : j \neq k \end{cases}. \quad (3.12)$$

Importantly, the covariance of the tests statistics does not depend on what model parameter is being tested, so a single null distribution can be used for tests of all parameters provided the tests are on the same DOF. This conserves computing time relative to the permutation procedure which must estimate a null distribution for each test.

3.3. The parametric-bootstrap procedure

We compute p -values using a parametric bootstrap: we use the estimate of $\hat{\Sigma}$ to generate B diagonal singular Wishart statistics. Because the rank of $\hat{\Sigma}$ is at most $\min\{(n - m), V\}$ it does not require the storage of the full $V \times V$ covariance matrix if $V > (n - m)$. This gives the following procedure for estimating the null distribution using the parametric bootstrap.

PROCEDURE 3.1 (Parametric bootstrap null estimation) Assuming the model (2.1):

1. Regress Y_v onto X to obtain the test statistics for $H_{v,0} : \beta_v = \beta_{v,0}$ using (3.9). Let $Z_{(1)0} < Z_{(2)0} < \dots < Z_{(V)0}$ denote the ascending test statistics and $\tilde{E} = [\hat{\epsilon}_{(1),0}, \dots, \hat{\epsilon}_{(V)0}]$, the $n \times V$ matrix of their associated residuals from model (2.1).
2. Standardize \tilde{E} so that the column norms are 1. Denote the standardized matrix by E . Let $r = \text{rank}(E) = \min\{n - m, V\}$.
3. Use E , m_1 , and r to generate the null test statistics $Z_b = (Z_{(1)b}, \dots, Z_{(V)b})$ for $b = 1, \dots, B$.
 - (a) Perform the singular value decomposition of $E = UDM^T$ where D is an $r \times r$ diagonal matrix and U and M have orthonormal columns. Let $\tilde{M} = \tilde{M}D$.
 - (b) Generate an $r \times m_1$ matrix S_b with independent standard normal entries.
 - (c) Obtain the null statistics $Z_b = \text{diag}(MS_b S_b^T M^T)$. Z_b are distributed according to a diagonal singular Wishart distribution.

The singular value decomposition only needs to be performed once for the entire procedure. Computing the statistics does not require multiplying $MS_b S_b^T M^T$, because we only need the diagonal entries. Procedures 2.3 and 2.4 are used to compute single-step and step-down adjusted p -values from the bootstrap sample.

4. METHODS

Code to perform the analyses presented in this manuscript is available at <https://bitbucket.org/simonvandekar/param-boot>. While the processed data are not publicly available, unprocessed data are available for download through the Database of Genotypes and Phenotypes (dbGaP) (Satterthwaite and others, 2016). We provide simulated region-wise data with the code so that readers can perform the region-wise analyses presented here. Synthetic simulations are used to assess finite sample properties of each testing procedure. Simulations that resample from the CBF data are used to assess the FWER in real data. Methods describing the simulation procedures are given in Section S1.2 of supplementary material available at Biostatistics online.

4.1. Cerebral blood flow data description

The Philadelphia Neurodevelopmental Cohort (PNC) is a large initiative to understand how brain maturation mediates cognitive development and vulnerability to psychiatric illness (Satterthwaite and others, 2014). In this study, we investigate image and regional measurements of CBF using an arterial spin labeling (ASL) sequence collected on 1,578 subjects, ages 8–21, from the PNC (see Satterthwaite and others, 2014 for details). Abnormalities in CBF have been documented in several psychiatric disorders including schizophrenia (Pinkham and others, 2011) and mood and anxiety disorders (Kaczkurkin and others, 2016). Establishing normative trajectories in CBF is critical to characterizing neurodevelopmental psychopathology (Satterthwaite and others, 2014).

Image preprocessing steps are described in Section S1.1 of supplementary material available at Biostatistics online. In order to parcellate the brain into anatomically defined regions, we use an advanced multi-atlas labeling approach which creates an anatomically labeled image in subject space (Avants and others, 2011). The ASL data are pre-processed using standard tools included with FSL (Jenkinson and others, 2002; Satterthwaite and others, 2014). The T1 image is then co-registered to the CBF image using boundary-based registration (Greve and Fischl, 2009), the transformation is applied to the label image, and then CBF values are averaged within each anatomically defined parcel.

Of the 1601 subjects with imaging data 23 did not have CBF data. Three hundred thirty-two were excluded due to clinical criteria, which include a history of medical disorders that affect the brain, a history of inpatient psychiatric hospitalization, or current use of a psychotropic medication. An additional 274 subjects were excluded by an imaging data quality assurance procedure, which included automatic and manual assessments of data quality and removal of subjects with negative mean CBF values in any of the anatomical parcels. These exclusions yielded a total of 972 subjects used for the imaging simulations and analysis.

4.2. Cerebral blood flow statistical analysis

We perform region- and voxel-wise analyses of CBF in order to identify locations where there are sex differences in development-related changes of CBF. For the region-wise analysis, we test the sex by age interaction on the average CBF trajectories in the $V = 112$ regions using an F-statistic from an unpenalized spline model. For the voxel-wise analysis, we perform the same test in all $V = 127\,756$ gray matter voxels.

For the region-wise data we fit the age terms with thin plate splines with 10 DOF. Thus, the numerator of the F-statistic has 9 DOF using the `mgcv` package in R (R Core Team, 2016; Wood, 2011). Results from the simulation analyses and previous theoretical results (Gotze, 1991) demonstrate that multivariate convergence rates depends on the dimension of the vector of statistics (Table 1 and Table S2 in supplementary material available at Biostatistics online), and the DOF of the test (Figure 2). For this reason, we use the Yeo-Johnson transformation for the PBJ procedure so that the transformed CBF data are approximately normal (Yeo and Johnson, 2000). We estimate the age terms for the voxelwise data on 5 DOF,

Table 1. *Type 1 error results for $n = 100$ to assess convergence rates. Values are mean percentage of correctly rejected tests across 500 simulations. Test statistics were simulated as normal with independent (Indep), positive autoregressive (Pos AR(1)), and negative autoregressive (Neg AR(1)) correlation structures with $\rho = 0.9$ and $\rho = -0.9$. The number of tests was varied within $m = (200, 500, 1,000, 5,000, 10,000)$ with 10% non-null test statistics. Detailed descriptions of the column names are given in Section 4.*

	$n = 100$	$T_n \mid \text{Holm}$	$Z_n \mid \text{Holm}$	$T_n \mid \Sigma$	$T_n \mid \hat{\Sigma}$	$Z_n \mid \hat{\Sigma}$	$T_n \mid \text{Perm}$
Indep	$m = 100$	6	6	6	8	6	4
	$m = 200$	6	4	7	8	5	5
	$m = 500$	7	6	8	8	6	5
	$m = 1000$	10	4	10	10	5	6
	$m = 5000$	10	4	11	10	4	5
	$m = 10\,000$	14	4	15	14	4	5
Pos AR(1)	$m = 100$	2	1	6	4	4	3
	$m = 200$	4	3	6	8	5	4
	$m = 500$	5	3	7	8	4	5
	$m = 1000$	7	2	8	9	5	6
	$m = 5000$	6	3	9	8	4	4
	$m = 10\,000$	9	3	11	10	4	3
Neg AR(1)	$m = 100$	5	3	7	9	6	5
	$m = 200$	3	2	5	6	4	3
	$m = 500$	4	2	8	8	5	4
	$m = 1000$	5	3	8	8	4	4
	$m = 5000$	8	4	10	9	6	5
	$m = 10\,000$	8	3	11	9	6	5

so that the test for the interaction is on 4 DOF. Race and motion (mean relative displacement; MRD) are included as covariates at each location for both analyses. Similar results were presented in a previous report (Satterthwaite and others, 2014) using Bonferroni adjustment. We also perform the voxel-wise analyses using a 10 DOF spline basis in Supplementary Material available at Biostatistics online.

The same MTPs are compared for the region- and voxel-wise CBF analysis as used in simulations. For the region- and voxel-wise analyses 10,000 samples are used for the PBJ and PJ procedures. We present the corrected results with the FWER controlled at $\alpha = 0.01$ for the region-wise analysis and $\alpha = 0.05$ for the voxel-wise analysis. A more conservative threshold is used for the region-wise data as the smaller number of comparisons and noise reduction from averaging within regions increases power considerably.

5. SIMULATION RESULTS

5.1. Synthetic simulation results

We use simulations to explore how the FWER is affected by using asymptotic approximations and estimating the covariance matrix in finite sample sizes. Results are shown for sample sizes of $n = 100$ and $n = 40$ (Table 1 and Table S2). Two hypotheses are tested on (see S1 of supplementary material available at Biostatistics online) and 3 (see S2 of supplementary material available at Biostatistics online) DOF. The results demonstrate that by using the sample covariance estimate in the PBJ MTP the FWER is only slightly inflated for $n = 40$ and when $n = 100$ the FWER is controlled at the nominal level for all the dimensions

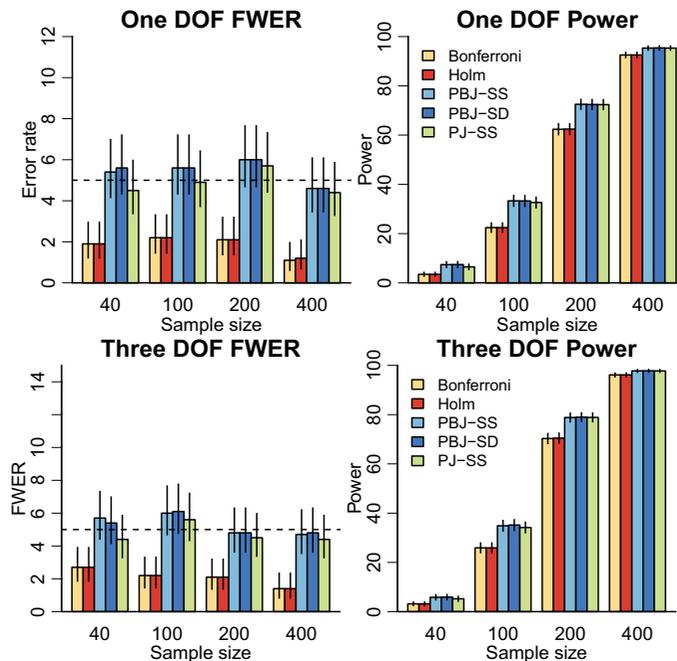


Fig. 1. FWER and power for the F-statistic of (see S1 of supplementary material available at Biostatistics online) on DOF and F-statistic of (see S2 of supplementary material available at Biostatistics online) on 3 DOF for $V = 112$ brain regions. Bonferroni and Holm both have conservative control. The PBJ maintains accurate FWE control even when $n < p$. The power for the PBJ and PJ procedures are equal. Lines indicate 95% confidence intervals. SS, single-step; SD, step-down. Color figure is available in the online version of this manuscript.

considered (Column $Z_n \mid \hat{\Sigma}$). Note, that when the transformation (3.9) is not used all procedures have inflated FWERs (Columns T_n in Table 1 and Table S2 of supplementary material available at Biostatistics online) due to the fact that the multivariate normal approximation is not accurate and is worse for a larger number of tests (Gotze, 1991). We can conclude that most of the error in estimating the FWER comes from the normal approximation. Even when $\hat{\Sigma}$ is rank deficient it still provides nominal FWER control. Interestingly, the PJ procedure controls the FWER for all the sample sizes considered. While permutation tests are exact for univariate distributions (Lehmann and Romano, 2005), to our knowledge there is no theoretical justification that multivariate permutations are accurate when the number of statistics exceeds the sample size. Finally, the PBJ and PJ MTPs control the FWER at the nominal level, while Holm's procedure is conservative for correlated test statistics.

5.2. Region-wise FWER and power

We use simulations that sample from real imaging data to assess the FWER in finite samples for region-wise analyses. For the test of hypothesis (see S1 of supplementary material available at Biostatistics online), the PBJ and PJ procedures maintain the nominal FWER for all samples (Figure 1). As expected, the FWER for the Bonferroni and Holm procedure are conservative. The power of the joint testing procedures is higher than the marginal testing procedures (Figure 1). The PBJ has equal power to the PJ procedure. The step-down procedure confers almost no benefit.

For testing the hypothesis (see S2 of supplementary material available at Biostatistics online) the PBJ and PJ procedures control the FWER at the nominal level for all sample sizes (Figure 1). The Bonferroni

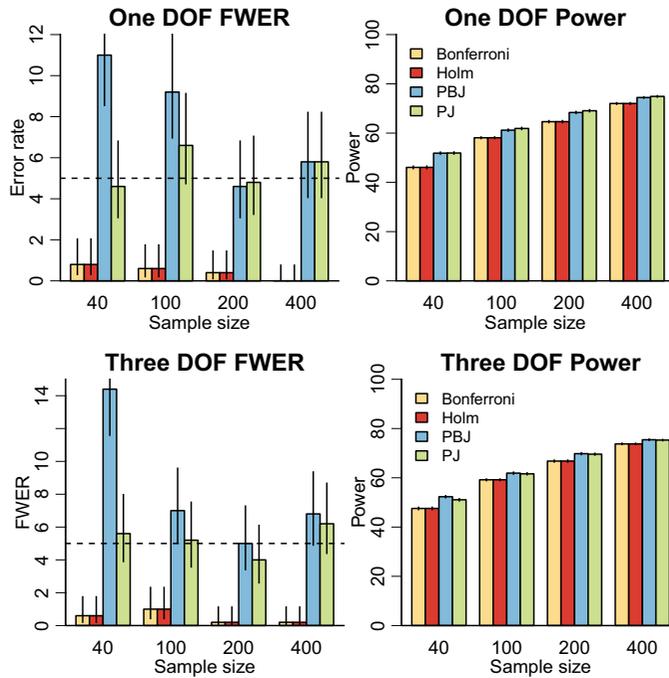


Fig. 2. FWER and power for the F-statistic of (see S1 of supplementary material available at Biostatistics online) on 1 DOF and F-statistic of (see S2 of supplementary material available at Biostatistics online) on 3 DOF for the CBF image with $V = 127\,756$ voxels. Bonferroni and Holm both have conservative control. The PBJ controls the FWER for samples sizes greater than 200. The power for the PBJ is approximately equal to the PJ tests. Lines indicate 95% confidence intervals. Color figure is available in the online version of this manuscript.

and Holm procedures give similarly conservative FWERs as the single DOF test. Power analyses demonstrate that the PBJ and PJ procedures have the same power. As with testing (see S1 of supplementary material available at Biostatistics online), the step-down procedure shows little improvement over the single-step.

5.3. Voxel-wise FWER and power

We use simulations that sample from real imaging data to assess the type 1 error and power for voxel-wise analyses. For the test of (see S1 of supplementary material available at Biostatistics online), the PBJ maintains the nominal FWER for samples sizes greater than 200. The PJ procedure maintains control for all sample sizes considered (Figure 2). As expected, Bonferroni and Holm procedures have conservative FWER. This conservative FWER leads to a reduction in power for these methods (Figure 2). The power for PBJ was comparable to the PJ procedure.

For the test of (see S2 of supplementary material available at Biostatistics online) on 3 DOF the PBJ controls the FWER for sample sizes larger than 200 (Figure 2). The PBJ procedure does not control the FWER for smaller sample sizes because the estimate of the covariance structure (2.2) is not accurate enough to work as a plug-in estimator for the full rank covariance matrix. The PJ procedure maintains nominal control of the FWER for all sample sizes. The marginal testing procedures have the same conservative FWER control as above. Both joint testing procedures have greater power than the marginal procedures.

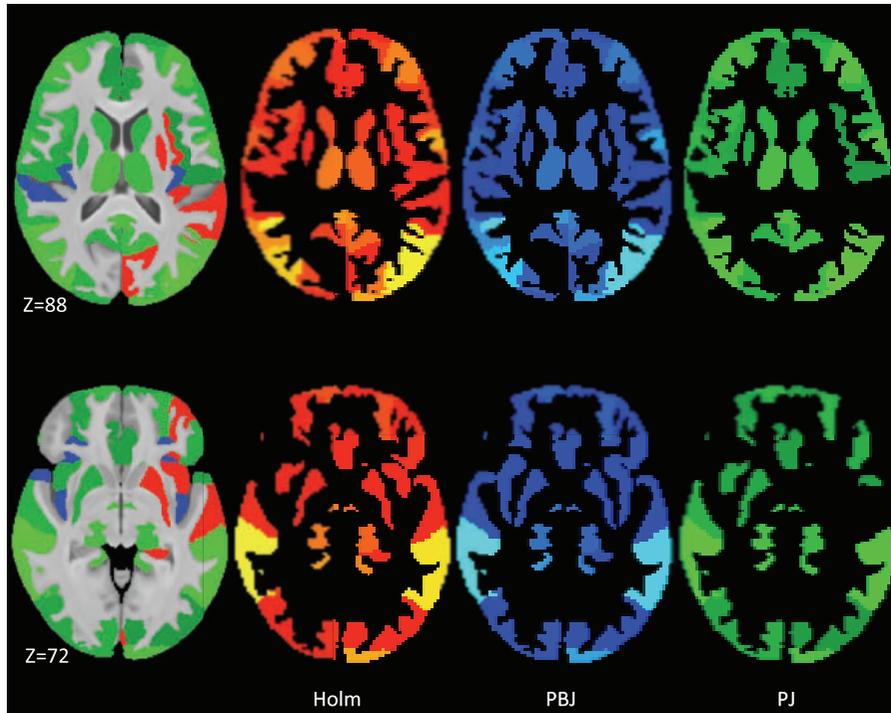


Fig. 3. FWER controlled results at $\alpha = 0.01$ for Holm, PBJ step-down, and PJ single-step for the region-wise analysis. Color scale is $-\log_{10}(p)$ and shows results greater than 2. The left-most images show the overlay of PBJ, Holm, and PJ in that order. The right three columns show regions identified by Holm, PBJ, and PJ, respectively. Color figure is available in the online version of this manuscript.

5.4. Computing time

While the PJ and PBJ procedures are both linear in the sample size, as the PBJ procedure works directly with the distribution of the test statistics we expect it to be faster than permutations. To compare observed computing times we took the ratio of the time to evaluate both tests for PJ and PBJ procedures. The computing times for the sample sizes simulated for the region-wise data are given in Figure S1 of supplementary material available at Biostatistics online. The PBJ is at least twice as fast as the PJ procedure and increases with the sample size so that the computing time reduction is larger for larger sample sizes.

6. CEREBRAL BLOOD FLOW RESULTS

To compare the testing procedures in the CBF data, we perform a test for a nonlinear age-by-sex interaction on CBF trajectories for region- and voxel-wise analyses. For the voxel-wise analysis the PBJ and PJ MTPs take 5.9 and 69.2 h to run, respectively. We perform an F-test for the interaction at each of the 112 regions (Figure 3). The Bonferroni, Holm, PBJ, and PJ procedures reject 56, 71, 78, and 61 regions, respectively. The PJ procedure is more conservative for two reasons. The first is that it is a single-step procedure; when there is a relatively large number of rejected tests then using a step-down procedure is more likely to improve power. The second reason is that the finite sample distribution of each of the regions is different. Regions near the edge of the brain are likely to be more heavily skewed due to imperfections in the image registration. By comparing all regions to the distribution of the maximum, the PJ procedure is necessarily

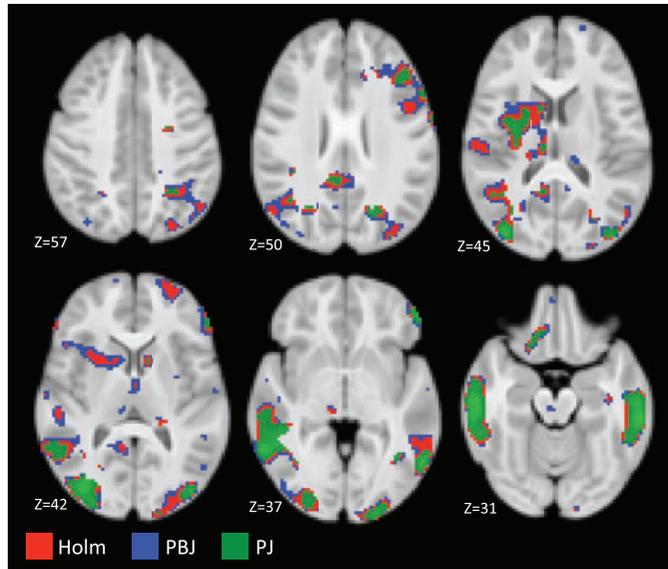


Fig. 4. FWER controlled results at $\alpha = 0.05$ for Holm, PBJ single-step, and PJ single-step for the voxel-wise analysis. Color scale is $-\log_{10}(p)$ for the adjusted p -values and shows results greater than 1.3. The overlay order is PBJ under Holm under PJ. Color figure is available in the online version of this manuscript.

conservative because it compares to the most heavily skewed regions. In contrast, by transforming the data prior to using the PBJ procedure the marginal distribution of the test statistics are approximately equal across regions.

For the voxel-wise analysis we perform the F-test for the nonlinear interaction on 4 DOF. The single-step PBJ offers improved power over the Holm and PJ procedures (Figure 4). The Holm procedure ignores the covariance structure of the test statistics so yields conservative results. The PJ procedure is more conservative even than the Holm procedure. As with the region-wise analysis, this is likely because the finite sample distribution of the test statistics is different: voxels near the edge of the brain tend to have higher variance and are likely heavily skewed. If there is a subset of voxels with a heavily skewed distribution, then taking the maximum test statistic will yield conservative inference for all locations that have tighter distribution. By transforming the distribution of the voxels to be approximately normal, the PBJ procedure offers improved power and speed.

7. DISCUSSION

We introduced a fast parametric bootstrap joint testing procedure as a new tool for multiple comparisons in neuroimaging. The PBJ procedure improves computing time by generating the test statistics directly instead of permuting the original data. If normality assumptions about the data generating distribution do not hold, then the Yeo–Johnson transformation can be used to obtain statistics that are approximately normal to improve the finite sample performance of the procedure.

In the CBF data analysis, the PBJ is more powerful than the PJ MTP because the PJ MTP does not account for the fact that the finite sample distribution of the test statistics can be different. Differences in the finite sample distribution of the statistics are attributable to certain regions near the edge of the brain having larger variance and skew. For this reason taking the maximum across locations leads to

conservative inference in locations that actually have tighter tails. While the PBJ generates from a chi-squared distribution this ensures that a few heavy-tailed locations do not affect the distribution of the maximum.

In simulations, the step-down procedures provide little improvement in power over the single-step procedures. However, in the regionwise analyses Holm rejected 15 more regions than the Bonferroni procedure. The reason for the difference is that step-down procedures offer little benefit when there is a small number of false null hypotheses and a large number of tests.

Using simulations we found that both joint procedures perform well, in some cases even when the number of tests exceeds the sample size. This is quite surprising as it seems impossible that any given estimate of the joint distribution will satisfactorily reproduce the true joint distribution of the test statistics. For example, if we consider the case of normal test statistics, $Z_n = (Z_{1n}, \dots, Z_{Vn}) \sim N(0, \Sigma)$ with full rank covariance, then the sufficient statistic is $\hat{\Sigma}_n$, which can be of rank n at most. So, the probabilities generated conditioning on $\hat{\Sigma}_n$ assume Z_n is restricted to a linear subspace of \mathbb{R}^V . With non-normal test statistics and more complex dependence structures it can only be more difficult to reproduce the null distribution.

MRI is a flexible noninvasive tool for studying neural aberrations related to psychiatric disorders such as schizophrenia (Pinkham and others, 2011) and mood and anxiety disorders (Kaczurkin and others, 2016). However, recent studies have shown that the PJ MTP is the only inference methodology to reliably control the FWER in neuroimaging data (Eklund and others, 2016; Silver and others, 2011). We have shown that inference using the currently available permutation procedure can take days and lead to conservative inference. Our proposed PBJ MTP is a reliable and fast testing procedure that will be a critical tool in studying functional and physiological features that can improve our understanding of the brain and its relation to behavior.

SUPPLEMENTARY MATERIAL

Supplementary material is available at <https://academic.oup.com/biostatistics>.

ACKNOWLEDGMENTS

We thank Mark van der Laan for a helpful discussion regarding null estimation.

FUNDING

RC2 grants from the National Institute of Mental Health MH089983 and MH089924. Center for Biomedical Computing and Image Analysis (CBICA) at Penn to R.T.S and T.D.S. National Institutes of Health grants: (T32MH065218 to S.N.V.); (R01MH107235 to R.C.G., R.E.G., R.T.S.), (R01MH107703 to T.D.S., R.T.S.), (R01MH112847 to TDS, KR, and RTS) and (R01NS085211 to R.T.S.). *Conflict of Interest:* None declared.

REFERENCES

- AVANTS, B. B., TUSTISON, N. J., SONG, G., COOK, P. A., KLEIN, A. AND GEE, J. C. (2011). A reproducible evaluation of ANTs similarity metric performance in brain image registration. *Neuroimage* **54**, 2033–2044.
- DAWID, A. P. (1981). Some matrix-variate distribution theory: notational considerations and a Bayesian application. *Biometrika* **68**, 265–274.
- DORAI-RAJ, S. (2014). *binom: Binomial Confidence Intervals For Several Parameterizations*. R package version 1.1-1.
- DUDOIT, S., SHAFFER, J. P. AND BOLDRICK, J. C. (2003). Multiple hypothesis testing in microarray experiments. *Statistical Science* **18**, 71–103.

- DUDOIT, S. AND VAN DER LAAN, M. J. (2008). *Multiple Testing Procedures with Applications to Genomics*, Springer Series in Statistics. New York, NY: Springer.
- DUNN, O. J. (1961). Multiple comparisons among means. *Journal of the American Statistical Association* **56**, 52–64.
- EKLUND, A., ANDERSSON, M., JOSEPHSON, C., JOHANNESSON, M. AND KNUTSSON, H. (2012). Does parametric fMRI analysis with SPM yield valid results? An empirical study of 1484 rest datasets. *NeuroImage* **61**, 565–578.
- EKLUND, A., NICHOLS, T. E. AND KNUTSSON, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences* **113**, 7900–7905.
- FREEDMAN, D. AND LANE, D. (1983). A nonstochastic interpretation of reported significance levels. *Journal of Business & Economic Statistics* **1**, 292–298.
- FRISTON, K. J., HOLMES, A. P., WORSLEY, K. J., POLINE, J.-P., FRITH, C. D. AND FRACKOWIAK, R. S. J. (1994a). Statistical parametric maps in functional imaging: a general linear approach. *Human Brain Mapping* **2**, 189–210.
- FRISTON, K. J., WORSLEY, K. J., FRACKOWIAK, R. S. J., MAZZIOTTA, J. C. AND EVANS, A. C. (1994b). Assessing the significance of focal activations using their spatial extent. *Human Brain Mapping* **1**, 210–220.
- GOTZE, F. (1991). On the rate of convergence in the multivariate CLT. *The Annals of Probability* **19**, 724–739.
- GREVE, D. N. AND FISCHL, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage* **48**, 63–72.
- HOCHBERG, Y. (1988). A sharper Bonferroni procedure for multiple tests of significance. *Biometrika* **75**, 800–802.
- HOLM, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics* **6**, 65–70.
- HOMMEL, G. (1988). A stagewise rejective multiple test procedure based on a modified Bonferroni test. *Biometrika* **75**, 383–386.
- JENKINSON, M., BANNISTER, P., BRADY, M. AND SMITH, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* **17**, 825–841.
- KACZKURKIN, A. N., MOORE, T. M., RUPAREL, K., CIRIC, R., CALKINS, M. E., SHINOHARA, R. T., ELLIOTT, M. A., HOPSON, R., ROALF, D. R., VANDEKAR, S. N. and others. (2016). Elevated amygdala perfusion mediates developmental sex differences in trait anxiety. *Biological Psychiatry* **80**, 775–785.
- KLEIN, A., ANDERSSON, J., ARDEKANI, B. A., ASHBURNER, J., AVANTS, B., CHIANG, M.-C., CHRISTENSEN, G. E., COLLINS, D. L., GEE, J., HELLIER, P. and others. (2009). Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *Neuroimage* **46**, 786–802.
- LEHMANN, E. L. AND ROMANO, J. P. (2005). *Testing Statistical Hypotheses*. New York, NY: Springer Science & Business Media.
- MARCUS, R., ERIC, P. AND GABRIEL, K. R. (1976). On closed testing procedures with special reference to ordered analysis of variance. *Biometrika* **63**, 655–660.
- PINKHAM, AMY, LOUGHEAD, JAMES, RUPAREL, KOSHA, WU, WEN-CHAU, OVERTON, EVE, GUR, RAQUEL AND GUR, RUBEN. (2011). Resting quantitative cerebral blood flow in schizophrenia measured by pulsed arterial spin labeling perfusion MRI. *Psychiatry Research* **194**, 64–72.
- R CORE TEAM. (2016). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- SARKAR, S. K. AND CHANG, C.-K. (1997). The Simes method for multiple hypothesis testing With positively dependent test statistics. *Journal of the American Statistical Association* **92**, 1601–1608.
- SATTERTHWAITE, T. D., CONNOLLY, J. J., RUPAREL, K., CALKINS, M. E., JACKSON, C., ELLIOTT, M. A., ROALF, D. R., HOPSON, R., PRABHAKARAN, K., BEHR, M. and others. (2016). The Philadelphia Neurodevelopmental Cohort: a publicly available resource for the study of normal and abnormal brain development in youth. *Neuroimage* **124**, 1115–1119.

- SATTERTHWAITE, T. D., ELLIOTT, M. A., RUPAREL, K., LOUGHEAD, J., PRABHAKARAN, K., CALKINS, M. E., HOPSON, R., JACKSON, C., KEEFE, J., RILEY, M., MENTCH, F. D., SLEIMAN, P., VERMA, R., DAVATZIKOS, C., HAKONARSON, H., GUR, R. C. *and others.* (2014a). Neuroimaging of the Philadelphia Neurodevelopmental Cohort. *NeuroImage* **86**, 544–553.
- SATTERTHWAITE, T. D., SHINOHARA, R. T., WOLF, D. H., HOPSON, R. D., ELLIOTT, M. A., VANDEKAR, S. N., RUPAREL, K., CALKINS, M. E., ROALF, D. R., GENNATAS, E. D. *and others.* (2014b). Impact of puberty on the evolution of cerebral perfusion during adolescence. *Proceedings of the National Academy of Sciences* **111**, 8643–8648.
- SILVER, M., MONTANA, G., NICHOLS, T. E., INITIATIVE, Alzheimer’s Disease Neuroimaging *and others.* (2011). False positives in neuroimaging genetics using voxel-based morphometry data. *Neuroimage* **54**, 992–1000.
- SRIVASTAVA, M. S. (2003). Singular Wishart and multivariate beta distributions. *The Annals of Statistics* **31**, 1537–1560.
- TUSTISON, N. J., AVANTS, B. B., COOK, P. A., ZHENG, Y., EGAN, A., YUSHKEVICH, P. A. AND GEE, J. C. (2010). N4itk: improved N3 bias correction. *IEEE Transactions on Medical Imaging* **29**, 1310–1320.
- Van der VAART, A. W. (2000). *Asymptotic Statistics*, Volume 3. New York, NY: Cambridge University Press.
- WESTFALL, P. H. AND YOUNG, S. S. (1993). *Resampling-based Multiple Testing: Examples and Methods for p-Value Adjustment*, Volume 279. New York, NY: John Wiley & Sons.
- WILSON, E. B. (1927). Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association* **22**, 209–212.
- WINKLER, A. M., RIDGWAY, G. R., WEBSTER, M. A., SMITH, S. M. AND NICHOLS, T. E. (2014). Permutation inference for the general linear model. *NeuroImage* **92**, 381–397.
- WOOD, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**, 3–36.
- YEO, I.-K. AND JOHNSON, R. A. (2000). A new family of power transformations to improve normality or symmetry. *Biometrika* **87**, 954–959.

[Received January 3, 2017; revised August 25, 2017; accepted for publication September 4, 2017]