

**Original citation:**

Gutierrez-Barragan, Felipe, Ithapu, Vamsi K., Hinrichs, Chris, Maumet, Camille, Johnson, Sterling C., Nichols, Thomas E. and Singh, Vikas. (2017) Accelerating permutation testing in voxel-wise analysis through subspace tracking : a new plugin for SnPM. NeuroImage, 159 . pp. 79-98.

**Permanent WRAP URL:**

<http://wrap.warwick.ac.uk/90817>

**Copyright and reuse:**

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

**Publisher's statement:**

© 2017, Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>

**A note on versions:**

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP url' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: [wrap@warwick.ac.uk](mailto:wrap@warwick.ac.uk)

# Accelerating Permutation Testing in Voxel-wise Analysis through Subspace Tracking: A new plugin for SnPM

Felipe Gutierrez-Barragan<sup>a</sup>, Vamsi K. Ithapu<sup>a</sup>, Chris Hinrichs<sup>a</sup>, Camille Maumet<sup>d</sup>,  
Sterling C. Johnson<sup>c</sup>, Thomas E. Nichols<sup>d</sup>, Vikas Singh<sup>b,a</sup>,  
for the Alzheimer’s Disease Neuroimaging Initiative <sup>1</sup>

<sup>a</sup>*Department of Computer Sciences, University of Wisconsin-Madison*

<sup>b</sup>*Department of Biostatistics & Med. Informatics, University of Wisconsin-Madison*

<sup>c</sup>*Department of Medicine, University of Wisconsin-Madison and William S. Middleton Veteran’s Hospital*

<sup>d</sup>*Department of Statistics, The University of Warwick*

<http://felipegb94.github.io/RapidPT/>

---

## Abstract

Permutation testing is a non-parametric method for obtaining the max null distribution used to compute corrected  $p$ -values that provide strong control of false positives. In neuroimaging, however, the computational burden of running such an algorithm can be significant. We find that by viewing the permutation testing procedure as the construction of a very large permutation testing matrix,  $\mathbf{T}$ , one can exploit structural properties derived from the data and the test statistics to reduce the runtime under certain conditions. In particular, we see that  $\mathbf{T}$  is low-rank plus a low-variance residual. This makes  $\mathbf{T}$  a good candidate for low-rank matrix completion, where only a very small number of entries of  $\mathbf{T}$  ( $\sim 0.35\%$  of all entries in our experiments) have to be computed to obtain a good estimate. Based on this observation, we present RapidPT, an algorithm that efficiently recovers the max null distribution commonly obtained through regular permutation testing in voxel-wise analysis. We present an extensive validation on a synthetic dataset and four varying sized datasets against two baselines: Statistical NonParametric Mapping (SnPM13) and a standard permutation testing implementation (referred as NaivePT). We find that RapidPT achieves its best runtime performance on medium sized datasets ( $50 \leq n \leq 200$ ), with speedups of 1.5x - 38x (vs. SnPM13) and 20x-1000x (vs. NaivePT). For larger datasets ( $n \geq 200$ ) RapidPT outperforms NaivePT (6x - 200x) on all datasets, and provides large speedups over SnPM13 when more than 10000 permutations (2x - 15x) are needed. The implementation is a standalone toolbox and also integrated within SnPM13, able to leverage multi-core architectures when available.

**Keywords:** Voxel-wise analysis, Hypothesis test, Permutation test, Matrix completion

---

Corresponding Author(s): Felipe Gutierrez Barragan (fgutierrez3@wisc.edu)

---

<sup>1</sup>Data used in preparation of this article was obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.ucla.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: [http://adni.loni.ucla.edu/wp-content/uploads/how\\_to\\_apply/ADNI\\_Acknowledgement\\_List.pdf](http://adni.loni.ucla.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf)

## 1. Introduction

Nonparametric voxel-wise analysis, e.g., via permutation tests, are widely used in the brain image analysis literature. Permutation tests are often utilized to control the family-wise error rate (FWER) in voxel-wise hypothesis testing. As opposed to parametric hypothesis testing schemes Friston et al. (1994); Worsley et al. (1992, 1996), nonparametric permutation tests Holmes et al. (1996); Nichols and Holmes (2002) can provide exact control of false positives while making minimal assumptions on the data. Further, despite the additional computational cost, permutation tests have been widely adopted in image analysis Arndt et al. (1996); Halber et al. (1997); Holmes et al. (1996); Nichols and Holmes (2002); Nichols and Hayasaka (2003) via implementations in broadly used software libraries available in the community SnPM (2013); FSL (2012); Winkler et al. (2014).

*Running time aspects of Permutation Testing.* Despite the varied advantages of permutation tests, there is a general consensus that the computational cost of performing permutation tests in neuroimaging analysis can often be quite high. As we will describe in more detail shortly, high dimensional imaging datasets essentially mean that *for each permutation*, hundreds of thousands of test statistics need to be computed. Further, as imaging technologies continue to get better (leading to higher resolution imaging data) and the concurrent slowdown in the predicted increase of processor speeds (Moore’s law), it is reasonable to assume that the associated runtime will continue to be a problem in the short to medium term. To alleviate these runtime costs, ideas that rely on code optimization and parallel computing have been explored Eklund et al. (2011); A. Eklund (2012, 2013). These are interesting strategies but any hardware-based approach will be limited by the amount of resources at hand. Clearly, significant gains may be possible if *more efficient schemes* that exploit the underlying structure of the imaging data were available. It seems likely that such algorithms can better exploit the resources (e.g., cloud or compute cluster) one has available as part of a study and may also gain from hardware/code improvements that are being reported in the literature.

Data acquired in many scientific studies, especially imaging and genomic data, are highly structured. Individual genes and/or individual voxels share a great deal of commonality with other genes and voxels. It seems reasonable that such correlation can be exploited towards better (or more efficient) statistical algorithms. For example, in genomics, Cheverud (2001) and Li and Ji (2005) used correlations in the data to estimate the effective number of independent tests in a genome sequence to appropriately threshold the test statistics. Also motivated by bioinformatics problems, Knijnenburg et al. (2009) approached the question of estimating the tail of the distribution of permutation values via an approximation by a generalized Pareto distribution (using fewer permutations). In the context of more general statistical analysis, the authors in Subramanian et al. (2005) proposed Gene Set Enrichment Analysis (GSEA) which exploits the underlying structure among the genes, to analyze gene-sets (e.g., where sets were obtained from biological pathways) instead of individual genes. If the genes within a gene-set have similar expression pattern, one may see improvements in statistical power. This idea of exploiting the “structure” towards efficiency (broadly construed) was more rigorously studied in Efron and Tibshirani (2007) and a nice non-parametric Bayesian perspective was offered in Dahl and Newton (2007). Within neuroimaging, a similar intuition drives Random Field theory based analysis Taylor and Worsley

(2008), albeit the objective there is to obtain a less conservative correction, rather than computational efficiency. Recently, motivated by neuroimaging applications and computational issues, Gaonkar and Davatzikos (2013) derived an analytical approximation of statistical significance maps to reduce the computational burden imposed by permutation tests commonly used to identify which brain regions contribute to a Support Vector Machines (SVM) model. In summary, exploiting the structure of the data to obtain alternative efficient solutions is not new, but we find that in the context of permutation testing on *imaging data*, there is a great deal of voxel-to-voxel correlations that if leveraged properly can, in principle, yield interesting new algorithms.

For permutation testing tasks in neuroimaging in particular, several groups have recently investigated ideas to make use of the underlying structure of the data to accelerate the procedure. In a preliminary conference paper (Hinrichs et al. (2013)), we introduced the notion of exploiting correlations in neuroimaging data via the underlying low-rank structure of the permutation testing procedure. A few years later, Winkler et al. (2016) presented the first thorough evaluation of the accuracy and runtime gains of six approaches that leverage the problem structure to accelerate permutation testing for neuroimage analysis. Among these approaches Winkler et al. (2016) presented an algorithm which relied on some of the ideas introduced by Hinrichs et al. (2013) to accelerate permutation testing through low-rank matrix completion (LRMC). Overall, algorithms that exploit the underlying structure of permutation testing in neuroimaging have provided substantial computational speedups.

### 1.1. Main Idea and Contributions

The starting point of our formulation is to analyze the entire permutation testing procedure via numerical linear algebra. In particular, the object of interest is the permutation testing matrix,  $\mathbf{T}$ . Each row of  $\mathbf{T}$  corresponds to the voxel-wise statistics, and each column is a specific permutation of the labels of the data. This perspective is not commonly used because a typical permutation test in neuroimaging rarely instantiates or operates on this matrix of statistics. Apart from the fact that  $\mathbf{T}$ , in neuroimaging, contains millions of entries, the reason for not working *directly* with it is because the goal is to derive the maximum null distribution. The central aspect of this work is to *exploit* the structure in  $\mathbf{T}$  – the spatial correlation across different voxel-statistics. Such correlations are not atypical because the statistics are computed from anatomically correlated regions in the brain. Even far apart voxel neighbourhoods are inherently correlated because of the underlying biological structures. This idea drives the proposed novel permutation testing procedure. We describe the contributions of this paper based on the observation that the permutation testing matrix is filled with related entries.

- **Theoretical Guarantees.** The basic premise of this paper is that *permutation testing in high-dimensions (especially, imaging data) is extremely redundant*. We show how we can model  $\mathbf{T}$  as a low-rank plus a low-variance residual. We provide two theorems that support this claim and demonstrate its practical implications. Our first result justifies this modeling assumption and several of the components involved in recovering  $\mathbf{T}$ . The second result shows that the error in the *global* maximum null distribution obtained from the estimate of  $\mathbf{T}$  is quite small.

- **A novel, fast and robust, multiple-hypothesis testing procedure.** Building upon the theoretical development, we propose a fast and accurate algorithm for permutation testing involving high-dimensional imaging data. The algorithm achieves state of the art runtime performance by estimating (or recovering) the statistics in  $\mathbf{T}$  rather than “explicitly” computing them. We refer to the algorithm as *RapidPT*, and we show that compared to existing state-of-the-art libraries for non-parametric testing, the proposed model achieves approximately  $20\times$  speed up over existing procedures. We further identify regimes where the speed up is even higher. *RapidPT* also is able to leverage serial and parallel computing environments seamlessly.
- **A plugin in SnPM (with stand-alone libraries).** Given the importance and the wide usage of permutation testing in neuroimaging (and other studies involving high-dimensional and multimodal data), we introduce a heavily tested implementation of *RapidPT* integrated as a plugin within the current development version of SnPM — a widely used non-parametric testing toolbox. Users can invoke *RapidPT* directly from within the SnPM graphical user interface and benefit from SnPM’s familiar pre-processing and post-processing capabilities. This final contribution, without a separate installation, brings the performance promised by the theorems to the neuroimaging community. Our documentation Gutierrez-Barragan and Ithapu (2016) gives an overview of how to use *RapidPT* within SnPM.

Although the present work shares some of the goals and motivation of Winkler et al. (2016) – specifically, utilizing the algebraic structure of  $\mathbf{T}$  – there are substantial technical differences in the present approach, which we outline further below. First, unlike Winkler et al. (2016), we directly study permutation testing for images at a more fundamental level and seek to characterize mathematical properties of relabeling (i.e., permutation) procedures operating on high-dimensional imaging data. This is different from assessing whether the underlying operations of classical statistical testing procedures can be reformulated (based on the correlations) to reduce computational burden as in Winkler et al. (2016). Second, by exploiting celebrated technical results in random matrix theory, we provide theoretical guarantees for estimation and recovery of  $\mathbf{T}$ . Few such results were known. Note that empirically, our machinery avoids a large majority of the operations performed in Winkler et al. (2016). Third, some speed-up strategies proposed in Winkler et al. (2016) can be considered as special cases of our proposed algorithm — interestingly, if we were to increase the number ‘actual’ operations performed by *RapidPT* (from  $\approx 1\%$ , suggested by our experiments, to  $50\%$ ), the computational workload begins approaching what is described in Winkler et al. (2016).

## 2. Permutation Testing in Neuroimaging

In this section, we first introduce some notations and basic concepts. Then, we give additional background on permutation testing for hypothesis testing in neuroimaging to motivate our formulation. Matrices and vectors will be denoted by bold upper-case and lower-case letters respectively, and scalars will be represented using non-bold letters. For a matrix  $\mathbf{X}$ ,  $\mathbf{X}[i, :]$  denotes the  $i^{th}$  row and  $\mathbf{X}[i, j]$  denotes the element in  $i^{th}$  row and  $j^{th}$  column.

Permutation testing is a nonparametric procedure for estimating the empirical distribution of the global null Edgington (1969b,a); Edgington and Onghena (2007). For a two-sample univariate statistical test, a permutation test computes an *unbiased* estimate of the null distribution of the relevant univariate statistic (e.g.,  $t$  or  $\chi^2$ ). Although univariate null distributions are, in general, well characterized, the sample maximum of the voxel-wise statistics usually does not have an analytical form due to strong correlations across voxels, as discussed in Section 1. Permutation testing is appropriate in this high-dimensional setting because it is non-parametric and does not impose any restriction on the correlations across the voxel-wise statistics. Indeed, when the test corresponds to group differences between samples based on a stratification variable, under the null hypothesis  $\mathcal{H}_0$ , the grouping labels given to the samples are artificial, i.e., they correspond to the set of all possible *relabellings* of the samples. In neuroimaging studies, typically the groups correspond to the presence or absence of an underlying disease condition (e.g., controls and diseased). Whenever  $\mathcal{H}_0$  is true, the data sample representing a healthy subject is, roughly speaking, ‘similar’ to a diseased subject. Under this setting, in principle, *interchanging* the labels of the two instances will have no effect on the distribution of the resulting voxel-wise statistics across all the dimensions (or covariates or features). So, if we randomly permute the labels of the data instances from both groups, under  $\mathcal{H}_0$  the recomputed sets of voxel-wise statistics correspond to the same global null distribution. We denote the number of such relabellings (or permutations) by  $L$ . The histogram of all  $L$  maximum statistics i.e., the maximum across all voxel-wise statistics for a given permutation, is the empirical estimate of the *exact* maximum null distribution under  $\mathcal{H}_0$ . When given the true/real labeling, to test for significant group-wise differences, one can simply compute the fraction of the max null that is more extreme than the maximum statistic computed across all voxels for this real labeling.

**The case for strong null.** Observe that when testing *multiple* sets of hypotheses there are two different types of control for the Type 1 error rate (FWER): weak and strong control Y. Hochberg (1987). A test is referred to as weak control for FWER whenever the Type 1 error rate is controlled only when all the hypotheses involved (here, the number of voxels being tested) are true. That is,  $\mathcal{H}_0$  is true for (all) voxels. On the other hand, a test provides strong control for FWER whenever Type 1 error rate is controlled under any combination/proportion of the true hypotheses. It is known that the procedure described above (i.e., using the max null distribution calculated across all voxel-wise statistics) provides strong control of FWER Holmes et al. (1996). This is easy to verify because the maximum of all voxel-wise statistics is used to compute the corrected  $p$ -value, and so, the exact proportion of which hypotheses are true does not matter. Further, testing based on strong control will classify non-activated voxels as activated with a probability upper bounded by  $\alpha$ , i.e., it has localizing power Holmes et al. (1996), a desirable property in neuroimaging studies in particular. For the remainder of this paper, we will focus on such strong control and restrict our presentation to the case of group difference analysis for two groups.

### 2.1. NAIVEPT: The exhaustive Permutation Testing procedure

Figure 1 and algorithm 1 illustrate the permutation testing procedure. Given the data instances from two groups,  $\mathbf{X}^1 \in \mathbb{R}^{v \times n_1}$  and  $\mathbf{X}^2 \in \mathbb{R}^{v \times n_2}$ , where  $n_1$  and  $n_2$  denote the number of subjects from each group respectively. Also,  $v$  denotes the number of voxels

in the brain image. Let  $n = n_1 + n_2$  and  $\mathbf{X} = [\mathbf{X}^1; \mathbf{X}^2]$  give the (row-wise) stacked data matrix ( $\mathbf{X} \in \mathbb{R}^{v \times n}$ ). Note that a permutation of the columns of  $\mathbf{X}$  corresponds to a group relabelling. The  $v$  distinct voxel-wise statistics are then computed for  $L$  such permutations, and used to construct the *permutation testing matrix*  $\mathbf{T} \in \mathbb{R}^{v \times L}$ . The empirical estimate of the max null is simply the histogram of the maximum of each of the columns of  $\mathbf{T}$  – denoted by  $h^L$ . Algorithm 1 is occasionally referred to as Monte Carlo permutation tests in the literature because of the random sampling involved in generating the statistics. This standard description of a permutation test will be used in the following sections to describe our proposed testing algorithm.

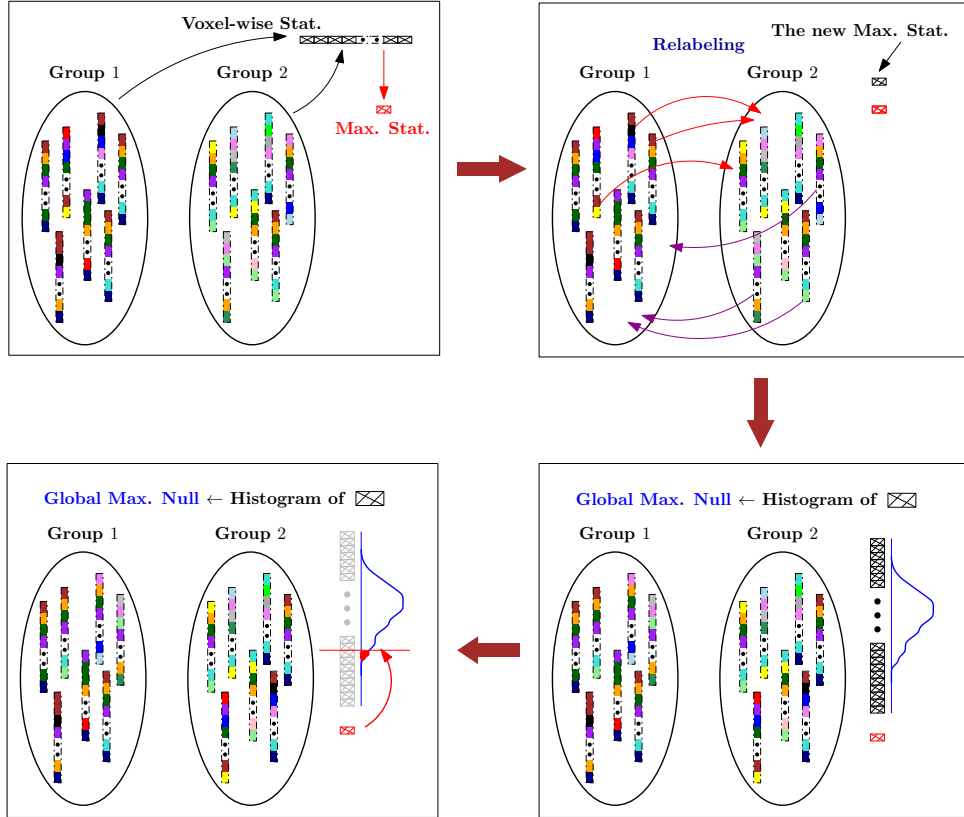


Figure 1: Flow diagram of the permutation testing procedure described in algorithm 1. Group 1 and Group 2 correspond to  $\mathbf{X}^1$  and  $\mathbf{X}^2$ , each new max stat corresponds to each  $m_i$ , and the global max null corresponds to  $h^L$ .

## 2.2. The situation when $L$ is large

Evidently as the number of permutations ( $L$ ) increases Algorithm 1 becomes extremely expensive. Nonetheless, performing a large number of permutations is essential in neuroimaging for various reasons. We discuss these reasons in some detail next.

- 1) Random sampling methods draw many samples from near the mode(s) of the distribution of interest, but fewer samples from the tail. To characterize the threshold for a small portion of the tail of a distribution, we must invariably draw a very large number of

---

**Algorithm 1** NAIVEPT The exhaustive permutation testing procedure.

---

**Input:**  $\mathbf{X}^1, \mathbf{X}^2, L$

**Output:**  $\mathbf{T}, \mathbf{h}^L$

$\mathbf{X} = [\mathbf{X}^1; \mathbf{X}^2], n = n_1 + n_2$

$m_1 \dots, m_L \leftarrow [\emptyset]$

**for**  $i \in 1, \dots, L$  **do**

$j_1 \dots, j_n \sim \text{PERMUTE}[1, n]$

$\tilde{\mathbf{X}}^1 \leftarrow \mathbf{X}[:, j_1, \dots, j_{n_1}], \tilde{\mathbf{X}}^2 \leftarrow \mathbf{X}[:, j_{n_1+1}, \dots, j_n]$

$\mathbf{T}[:, i] \leftarrow \text{test}(\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2)$

$m_i \leftarrow \text{MAX}(\mathbf{T}[:, i])$

**end for**

$\mathbf{h}^L \leftarrow \text{HISTOGRAM}(m_1, \dots, m_L)$

---

samples just so that the estimate converges. So, if we want an  $\alpha = 0.01$  threshold from the max null distribution, we require many thousands of permutation samples — a computationally expensive procedure in neuroimaging as previously discussed.

- 2) Ideally we want to obtain a precise threshold for any  $\alpha$ , in particular small  $\alpha$ . However, the smallest possible  $p$ -value that can be obtained from the empirical null is  $\frac{1}{L}$ . Therefore, to calculate very low  $p$ -values (essential in many applications),  $L$  must be very large.
- 3) A typical characteristic of brain imaging disorders, for instance in the *early* (e.g., preclinical) stages of Alzheimer’s disease (AD) and other forms of dementia, is that the disease signature is subtle — for instance, in AD, the deposition of Amyloid load estimated via positron emission tomographic (PET) images or atrophy captured in longitudinal structural magnetic resonance images (MRI) image scans in the asymptomatic stage of the disease. The signal is weak in this setting and requires large sample size studies (i.e., large  $n$ ) and a need for estimating the Type 1 error threshold with high confidence. The necessity for high confidence tail estimation implies that we need to sample many more relabelings, requiring a large number of permutations  $L$ .

### 3. A Convex Formulation to characterize the structure of $\mathbf{T}$

It turns out that the computational burden of algorithm 1 can be mitigated by exploiting the structural properties of the permutation testing matrix  $\mathbf{T}$ . Our strategy uses ideas from LRMC, subspace tracking, and random matrix theory, to exploit the correlated structure of  $\mathbf{T}$  and model it in an alternative form. In this section, we first introduce LRMC, followed by the overview of the eigen-spectrum properties of  $\mathbf{T}$ , which then leads to our proposed model of  $\mathbf{T}$ .

#### 3.1. Low-Rank Matrix Completion

Given only a small fraction of the entries in a matrix, the problem of low-rank matrix completion (LRMC) Candès and Tao (2010) seeks to *recover* the missing entries of the entire matrix. Clearly, with no assumption on the properties of the matrix, such a recovery is ill-posed. Instead, if we *assume* that the column space of the matrix is low-rank and the observed entries are randomly sampled, then the authors of Candès and Tao (2010)



and others have shown that, with sufficiently small number of entries, one can recover the orthogonal basis of the row space as well as the expansion coefficients for each column — that is, fully recover the missing entries of the matrix. Specifically, the number of entries required is roughly  $r \log(d)$  where  $r$  is the column space’s rank and  $d$  is the ambient dimension. By placing an  $\ell_1$ -norm penalty on the eigenvalues of the recovered matrix, i.e., the nuclear norm Fazel et al. (2004); Recht et al. (2010), one optimizes a convex relaxation of an (non-convex) objective function which explicitly minimizes the rank. Alternatively, we can specify a rank  $r$  ahead of time, and estimate an orthogonal basis of that rank by following a gradient along the Grassmannian manifold Balzano et al. (2010); He et al. (2012). The LRMC problem has received a great deal of attention in the period after the Netflix Prize Bennett and Lanning (2007), and numerous applications in machine learning and computer vision have been investigated Ji et al. (2010). Details regarding existing algorithms and their analyses including strong recovery guarantees are available in Candès and Recht (2009); Recht (2011).

**LRMC Formulation.** Let us consider a matrix  $\mathbf{T} \in \mathbb{R}^{v \times L}$ . Denote the set of randomly subsampled entries of this matrix as  $\Omega$ . This means that we have access to  $\mathbf{T}_\Omega$ , and our recovery task corresponds to estimating  $\mathbf{T}_{\Omega^C}$ , where  $\Omega^C$  corresponds to the complement of the set  $\Omega$ . Let us denote the estimate of the complete matrix be  $\hat{\mathbf{T}}$ . The completion problem can be written as the following optimization task,

$$\min_{\hat{\mathbf{T}}} \|\mathbf{T}_\Omega - \hat{\mathbf{T}}_\Omega\|_{\text{Frob}}^2 \quad (1)$$

$$\text{s.t. } \hat{\mathbf{T}} = \mathbf{U}\mathbf{W} \quad (2)$$

$$\mathbf{U} \text{ is orthogonal.} \quad (3)$$

where  $\mathbf{U} \in \mathbb{R}^{v \times r}$  is the low-rank basis of  $\mathbf{T}$ , i.e., the columns of  $\mathbf{U}$  correspond to the orthogonal basis vectors of the column space of  $\mathbf{T}$ . Here,  $\Omega$  gives the measured entries and  $\mathbf{W}$  is the matrix of coefficients that lets us reconstruct  $\hat{\mathbf{T}}$ .

### 3.2. Low-rank plus a long tail in $T$

Most datasets encountered in the real world (and especially in neuroimaging) have a dominant low-rank component. While the data may not be *exactly* characterized by a low-rank basis, the residual will not significantly alter the eigen-spectrum of the sample covariance in general. Strong correlations nearly always imply a skewed eigen-spectrum, because as the the eigen-spectrum becomes flat, the resulting covariance matrix tends to become sparser (the “uncertainty principle” between low-rank and sparse matrices Chandrasekaran et al. (2011)). Low-rank structure in the data is encountered even more frequently in neuroimaging — unlike natural images in computer vision, there is much stronger voxel-to-voxel homogeneity in a brain image.

While performing statistical hypothesis testing on these images, the low-rank structure described above carries through to  $\mathbf{T}$  for purely linear statistics such as sample means, mean differences and so on. However, non-linearities in the test statistic, e.g., normalizing by pooled variances, will perturb the eigen-spectrum of the original data, contributing a long tail of eigenvalues (see Figure 2). This large number of significant singular values needs to be accounted for, if one intends to model  $\mathbf{T}$  using low-rank structure. Ideally, we require that this long tail should either decay rapidly, or that it does not overlap with the dominant

eigenvalues. This is equivalent to asking that the resulting non-linearities do not *decorrelate* the test statistics, to the point that the matrix  $\mathbf{T}$  cannot be approximated by a low-rank matrix with high fidelity. For  $t$ -statistics, the non-linearities come from normalization by pooled variances, see for example a two-sample  $t$ -test shown in (4). Here  $(\mu_1, \sigma_1)$  and  $(\mu_2, \sigma_2)$  are the mean and standard deviations for the two groups respectively. Since the pooled variances calculated from correlated data  $X$  are unlikely to change very much from one permutation sample to another (except outliers), we expect that the spectrum of  $\mathbf{T}$  will resemble that of the data (or sample) covariance, *with the addition of a long, exponentially decaying tail*. More generally, if the non-linearity does not decorrelate the test statistics too much, it will almost certainly preserve the low-rank structure.

$$t = \frac{\mu_1 - \mu_2}{\sqrt{\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}}} \quad (4)$$

### Does low-rank plus a long tail assumption hold for other image modalities?

The underlying thesis of our proposed framework is that the permutation testing matrix  $\mathbf{T}$ , in general, has this low-rank plus long tail structure. In Figure 2, we show evidence that this is in fact the case for a variety of imaging modalities that are commonly used in medical studies – Arterial Spin Labeling (ASL), Magnetic Resonance Imaging (MRI), and two Positron Emission Tomography (PET) modalities including fluorodeoxyglucose (FDG) and Pittsburgh compound B (PiB). Using several images coming from each of these modalities four different  $\mathbf{T}$ s are constructed. Each row in Figure 2 shows the decay of the singular value spectrum of these four different  $\mathbf{T}$ s. The low-rank (left column) and the remaining long tail (right column) is clearly seen in these spectrum plots which suggests that the core modeling assumptions are satisfied. We note that the MRI data that was used in Figure 2 was in fact the de facto dataset for our evaluation presented in section 6. The underlying construct of such high correlations across multiple covariates or predictors is common to most biological datasets beyond brain scans, like genetic datasets.

### 3.3. Overview of Proposed method

If the low-rank structure dominates the long tail described above, then its contribution to  $\mathbf{T}$  can be modeled as a low variance Gaussian I.I.D. residual. A Central Limit argument appeals to the number of independent eigenfunctions that contribute to this residual, and, the orthogonality of eigenfunctions implies that as more of them meaningfully contribute to each entry in the residual, the more independent those entries become. In other words, if this long tail begins at a low magnitude and decays slowly, then we can treat it as a Gaussian I.I.D. residual; and if it decays rapidly, then the residual will perhaps be less Gaussian, but also more negligible. Thus, our algorithm makes no direct assumption about these eigenvalues themselves, but rather that the residual corresponds to a low-variance I.I.D. Gaussian random matrix – its contribution to the covariance of test statistics will be *Wishart* distributed, and from this property, we can characterize its eigenvalues.

**Why should we expect runtime improvements?** The low-rank + long tail structure of the permutation testing matrix then translates to the following identity,

$$\mathbf{T} = \mathbf{U}\mathbf{W} + \mathbf{S} \quad (5)$$

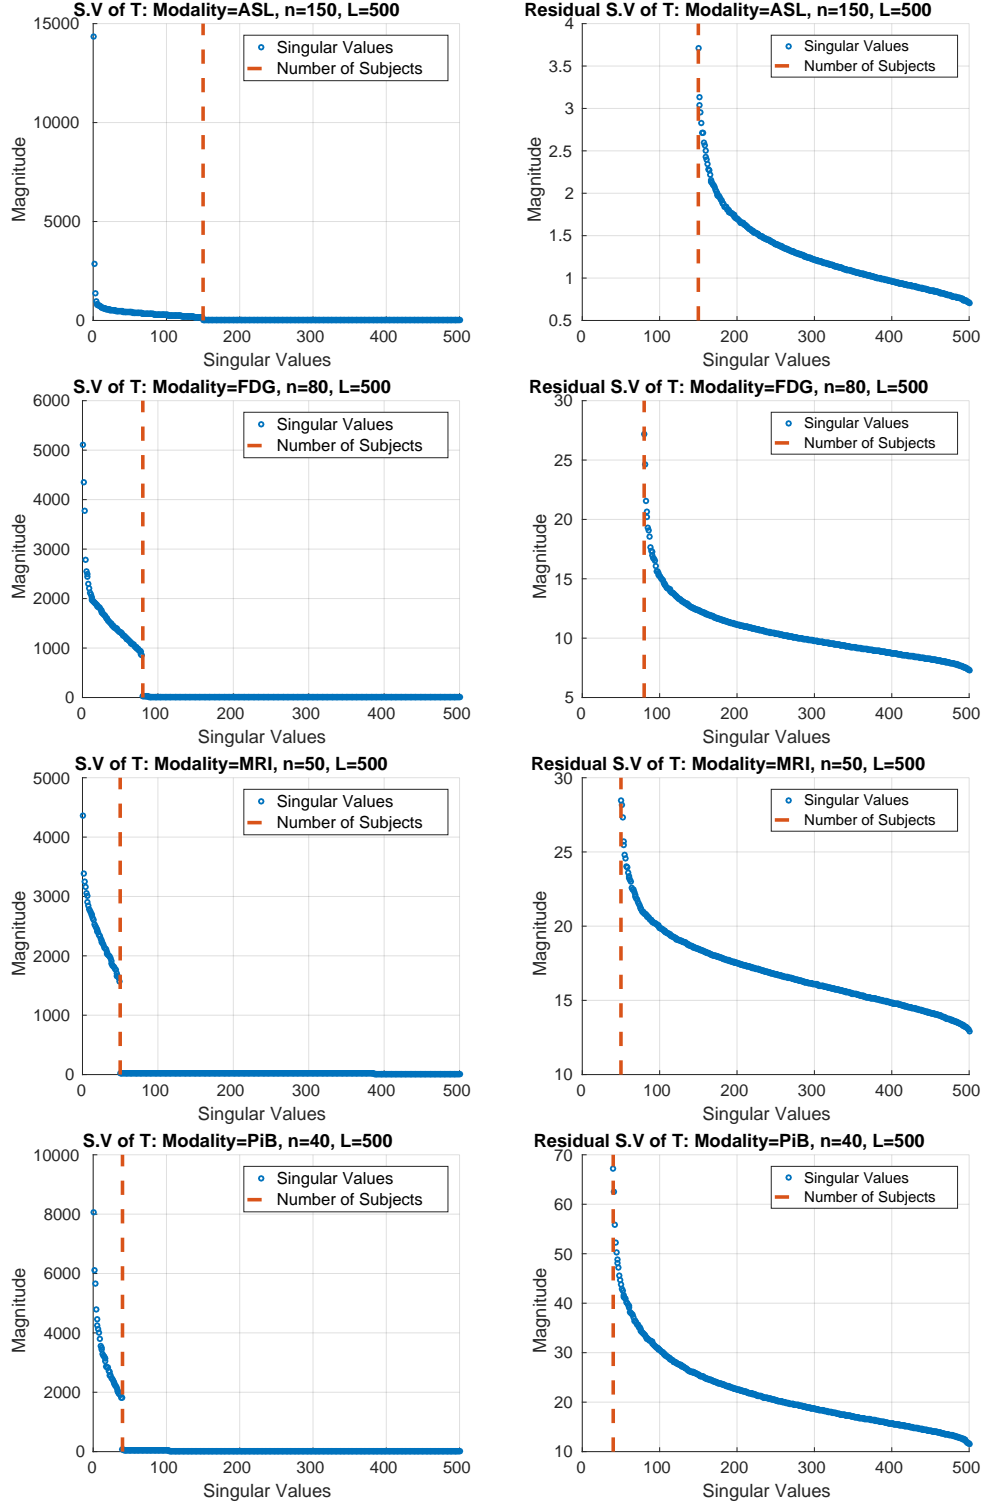


Figure 2: Singular value spectrums for permutation testing matrices with dimensions generated from the imaging modalities: ASL, FDG PET, MRI, and PiB PET. *Left:* Full spectrum of  $\mathbf{T}$  with  $L$  rows and  $v$  columns. *Right:* Residual singular values of  $\mathbf{T}$ .

where  $\mathbf{S}$  is the *unknown* I.I.D. Gaussian random matrix. We do not restrict ourselves to one-sided tests here, and so,  $\mathbf{S}$  is modeled to be zero-mean. Later in Section 4, we show that this apparent zero-mean assumption is addressed because of a post-processing step. The low-rank portion of  $\mathbf{T}$  can be reconstructed by sub-sampling the matrix at  $\Omega$  using the LRMC optimization from (1). Recall from the discussion in Section 3.1 that  $\Omega$  corresponds to a subset of indices of the entries in  $\mathbf{T}$ , i.e., instead of computing all voxel-wise statistics for a given relabeling (a column of  $\mathbf{T}$ ), only a small fraction  $\eta$ , referred to as the sub-sampling rate, are computed. Later in Sections 5 and 6, we will show that  $\eta$  is very small (on the orders of  $< 1\%$ ). Therefore, the overall number of entries in  $\Omega$  — *the number of statistics actually calculated to recover  $\mathbf{T}$*  — is  $\eta vL$  as opposed to  $vL$  for  $\eta \ll 1$ .

Since the core of the proposed method is to model  $\mathbf{T}$  by accessing only a small subset of its entries  $\Omega$ , we refer to it as a *rapid permutation testing* procedure – RAPIDPT. Observe that a large contributor to the running time of online subspace tracking algorithms, including the LRMC optimization from (1), is the module which updates the basis set  $\mathbf{U}$ ; but once a good estimate for  $\mathbf{U}$  has been found, this additional calculation is no longer needed. Second, the eventual goal of the testing procedure is to recover the max null as discussed earlier in Section 2, which then implies that the residual  $\mathbf{S}$  should also be recovered with high fidelity. Exact recovery of  $\mathbf{S}$  is not possible. Although, for our purposes, we only need its effect on the *distribution of the maximum* per permutation test. An estimate of the mean and variance of  $\mathbf{S}$  then provides reasonably good estimates of the max null. We therefore divide the entire process into two steps: *training*, and *recovery* which is described in detail in the next section.

## 4. Rapid Permutation Testing – RapidPT

In this section, we discuss the specifics of the training and recovery stages of RAPIDPT, and then present the complete algorithm, followed by some theoretical guarantees regarding consistency and recovery of  $\mathbf{T}$ . Figure 3 shows a graphical representation of the RapidPT algorithm 2.

### 4.1. The training phase

The goal of the training phase is to estimate the basis  $\mathbf{U}$ . Here, we perform a small number of fully sampled permutation tests, i.e., for approximately a few hundred of the columns of  $\mathbf{T}$  (each of which corresponds to a permutation) denoted by  $\ell$ , all the  $v$  voxel-wise statistics are computed. This  $v \times \ell$  “sub-matrix” of  $\mathbf{T}$  is referred to as the *training set*, denoted by  $\mathbf{T}_{ex}$ . In our experiments,  $\ell$  was selected to be either a fraction or a multiple of the total number of subjects  $n$  as described in section 5.4. From  $\mathbf{T}_{ex}$ , we estimate the basis  $\mathbf{U}$  using sub-sampled matrix completion methods Balzano et al. (2010); He et al. (2012), making *multiple* passes over the training set with the (given) sub-sampling rate  $\eta$ , until convergence. This corresponds to initializing  $\mathbf{U}$  as a random orthogonal matrix of a pre-determined rank  $r$ , and using the columns of  $\mathbf{T}_{ex}$  repeatedly to iteratively update it until convergence (see Balzano et al. (2010); He et al. (2012) for details regarding subspace tracking). Once  $\mathbf{U}$  is obtained in this manner,  $\mathbf{W}_{ex}$  is obtained by running a simple least-squares procedure on  $\mathbf{T}_{ex}$  and  $\mathbf{U}$ . The histogram of  $\mathbf{T}_{ex} - \mathbf{U}\mathbf{W}_{ex}$  will then be an estimate of the empirical distribution of the residual  $\mathbf{S}$  over the training set. We denote the standard deviation of these ‘left over’ entries as  $\sigma$ . We now discuss a few relevant aspects of this training phase.

Notice that in principle, one can estimate  $\mathbf{U}$  directly from  $\mathbf{T}_{ex}$  by simply computing the leading  $r$  principal components. This involves a brute-force approximation of  $\mathbf{U}$  by computing the singular-value decomposition of a dense  $v \times \ell$  matrix. Even for reasonably small  $v$ , this is a costly operation. Second,  $\hat{\mathbf{T}}$ , by definition, contains a non-trivial residual. We have no direct control on the structure of  $\mathbf{S}$  except that it is I.I.D Gaussian. Clearly, the variance of entries of  $\mathbf{S}$  will depend on the fidelity of the approximation provided by  $\mathbf{U}$ . Since the sub-sampling rate  $\eta$  (the size of the set  $\Omega$  compared to  $vL$ ) is known ahead of time, estimating  $\mathbf{U}$  via a subspace-tracking procedure using  $\eta$  fraction of the entries of  $\mathbf{T}_{ex}$  (where each column of  $\mathbf{T}_{ex}$  modifies an existing estimate of  $\mathbf{U}$ , one-by-one, without requiring to store all the entries of  $\mathbf{T}_{ex}$ ) directly provides an estimate of  $\mathbf{S}$ .

**Bias-Variance Trade-off.** When using a very sparse subsampling method i.e., sampling with small  $\eta$ , there is a bias-variance trade-off in estimating  $\mathbf{S}$ . Clearly, if we use the entire matrix  $\mathbf{T}$  to estimate  $\mathbf{U}$ ,  $\mathbf{W}$  and  $\mathbf{S}$ , we will obtain reliable estimates of  $\mathbf{S}$ . But, there is an overfitting problem: the least-squares objective used in fitting  $\mathbf{W}$  (in getting a good estimate of the max null) to such a small sample of entries is likely to grossly underestimate the variance of  $\mathbf{S}$  compared to when we use the entire matrix; the sub-sampling problem is not nearly as over-constrained as it is for the full matrix. This sampling artifact reduces the apparent variance of  $\mathbf{S}$ , and induces a bias in the distribution of the sample maximum, because extreme values are found less frequently. This sampling artifact has the effect of “shifting” the distribution of the sample maximum towards zero. We refer to this as a bias-variance trade-off because, we can think of the need for shift as an outcome of the sub-optimality of the estimate of  $\sigma$  versus the deviation of the true max null from the estimated max null. We correct for this bias by estimating the amount of the shift during the training phase, and then shifting the recovered sample max distribution by this estimated amount. This shift is denoted by  $\mu$ .

#### 4.2. The recovery phase

In the recovery phase, we sub-sample a small fraction of the entries of each column of  $\mathbf{T}$  successively, i.e., for each new relabeling, the voxel-wise statistics are computed over a small fraction of all the voxels to populate  $\mathbf{T}_\Omega$ . Using this  $\mathbf{T}_\Omega$ , and the pre-estimated  $\mathbf{U}$ , the reconstruction coefficients for this column  $\mathbf{w} \in \mathbb{R}^{r \times 1}$  are computed. After adding the random residuals – I.I.D Gaussian with mean  $\mu$  and standard deviation  $\sigma$ , to this  $\mathbf{U}\mathbf{w}$ , we have our estimate  $\hat{\mathbf{T}}$  for this specific relabeling/permutation. Recall that  $\mathbf{S}$  was originally modeled to be zero-mean (see (5)), but the presence of the shift  $\mu$  suggests a  $\mathcal{N}(\mu, \sigma^2)$  distribution instead. Overall, this entails recovering a total of  $v$  voxel-wise statistics from  $\eta v$  of such entries where  $\eta \ll 1$ . This process repeats for all the remaining  $L - \ell$  columns of  $\mathbf{T}$ , eventually providing  $\hat{\mathbf{T}}$ . Once  $\hat{\mathbf{T}}$  has been estimated, we proceed exactly as in the NAIVEPT, to compute the max null and test for the significance of the true labeling.

#### 4.3. The Algorithm

Algorithm 2 and Figure 3 summarizes RAPIDPT. The algorithm takes in the input data  $\mathbf{X}$ , the rank of the basis  $r$ , the sub-sampling rate  $\eta$ , the number of training columns  $\ell$  and the total number of columns  $L$  as inputs, and returns the estimated permutation testing matrix  $\mathbf{T}$  and the max null distribution  $h^L$ . As described earlier in Sections 4.1 and 4.2, Algorithm

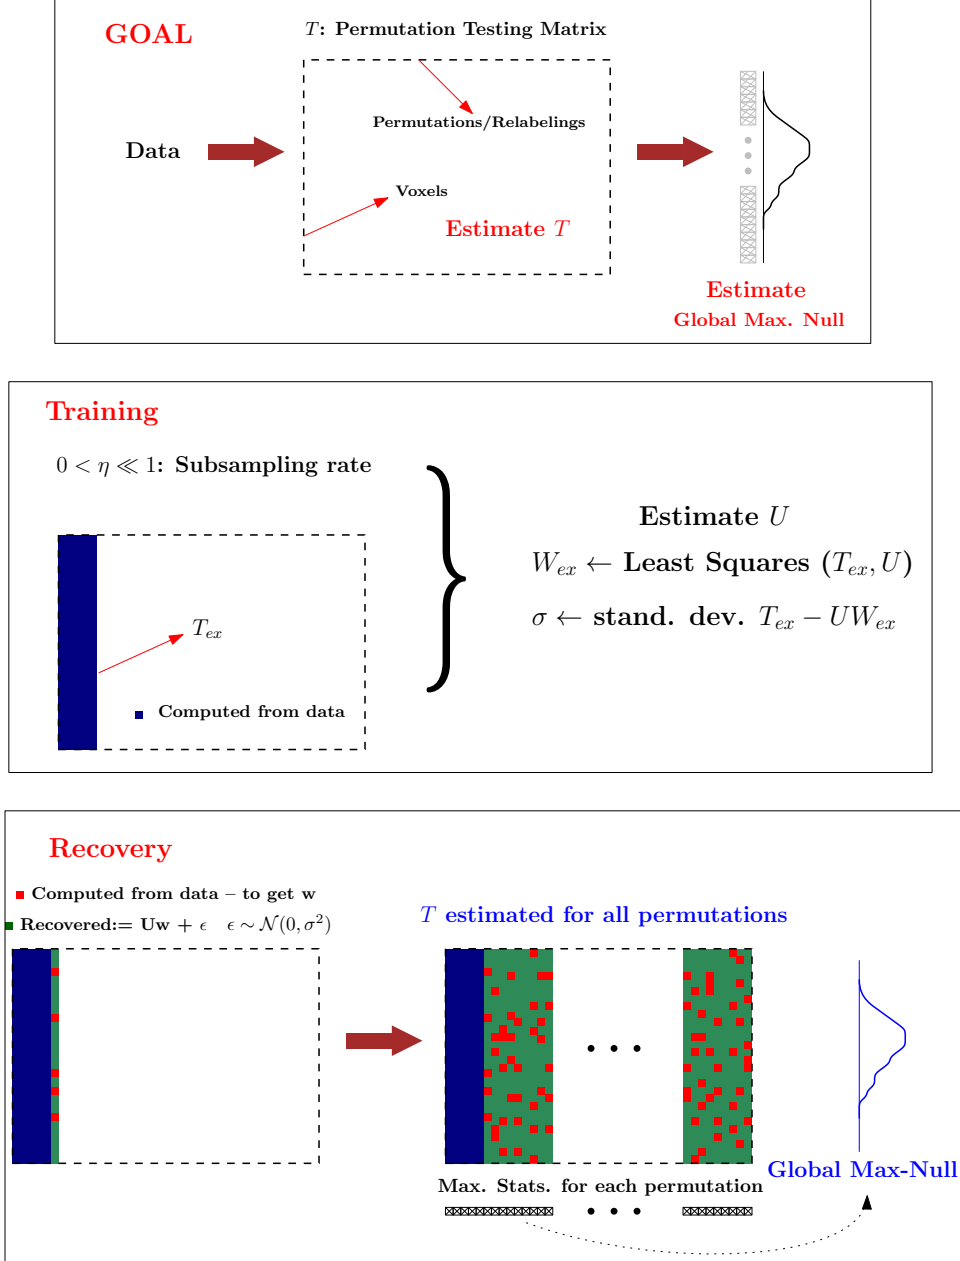


Figure 3: Flow diagram of the training and recovery steps in RapidPT. The number of training samples ( $l$ ) and the rank of  $\mathbf{U}(r)$  is the number of columns computed in the training phase (the blue area). The sub-sampling rate,  $\eta$ , is the fraction of red over green entries computed per column. The global max null is  $h^L$  in the algorithm.

2 first estimates  $\mathbf{U}$ ,  $\sigma$  and the shift  $\mu$ , which are then used to compute  $\mathbf{W}$  and  $\mathbf{S}$  for the  $L$  number of permutations.

#### 4.4. Summary of theoretical guarantees

Algorithm 2 shows an efficient way to recover the max null by modeling  $\mathbf{T}$  as a low-rank ( $\mathbf{UW}$ ) plus a low-variance residual ( $\mathbf{S}$ ). While this is useful, one may ask if this

---

**Algorithm 2** The RAPIDPT algorithm for permutation testing.

---

**Input:**  $\mathbf{X}^1, \mathbf{X}^2, r, \eta, L, \ell, \text{stat}$

**Output:**  $\hat{\mathbf{T}}, h^L$

$\mathbf{X} = [\mathbf{X}^1; \mathbf{X}^2], n = n_1 + n_2$

TRAINING

$\mathbf{U} \leftarrow \text{RAND. ORTH.}, \mathbf{W}_{ex} = [\emptyset]$

**for**  $i \in 1, \dots, \ell$  **do**

$j_1, \dots, j_n \sim \text{PERMUTE}[1, n]$

$\tilde{\mathbf{X}}^1 \leftarrow \mathbf{X}[:, j_1, \dots, j_{n_1}]$

$\tilde{\mathbf{X}}^2 \leftarrow \mathbf{X}[:, j_{n_1+1}, \dots, j_n]$

$\mathbf{T}_{ex}[:, i] \leftarrow \text{test}(\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2)$

$k_1, \dots, k_{\lceil \eta v \rceil} \sim \text{UNIF}[1, v]$

$\tilde{\mathbf{T}} \leftarrow \mathbf{T}_{ex}[k_1, \dots, k_{\lceil \eta v \rceil}, i]$

$\mathbf{U}, \mathbf{W}_{ex}[:, i] \leftarrow \text{SUBSPACE-TRACKING}(r)$

**end for**

$\sigma \leftarrow \text{STANDARD DEVIATION}\{\mathbf{T}_{ex} - \mathbf{U}\mathbf{W}_{ex}\}_{\Omega}$

$\mu \leftarrow \sup_i \text{MAX}\{\mathbf{T}_{ex}[:, i] - \mathbf{U}\mathbf{W}_{ex}[:, i]\}$

**for**  $i \in 1, \dots, \ell$  **do**

$\hat{\mathbf{T}}[:, i] \leftarrow \mathbf{T}[:, i]$

**end for**

RECOVERY

**for**  $i \in \ell + 1, \dots, L$  **do**

$k_1, \dots, k_{\lceil \eta v \rceil} \sim \text{UNIF}[1, v]$

$j_1, \dots, j_n \sim \text{PERMUTE}[1, n]$

$\tilde{\mathbf{X}}^1 \leftarrow \mathbf{X}[k_1, \dots, k_{\lceil \eta v \rceil}, j_1, \dots, j_{n_1}]$

$\tilde{\mathbf{X}}^2 \leftarrow \mathbf{X}[k_1, \dots, k_{\lceil \eta v \rceil}, j_{n_1+1}, \dots, j_n]$

$\tilde{\mathbf{T}} \leftarrow \text{test}(\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2)$

$\mathbf{W}[:, i] \leftarrow \text{COMPLETE}(\mathbf{U}, \tilde{\mathbf{T}}, k_1, \dots, k_{\lceil \eta v \rceil})$

$\mathbf{s} \leftarrow \text{i.i.d.}\mathcal{N}^v(0, \sigma^2)$

$\hat{\mathbf{T}}[:, i] \leftarrow \mathbf{U}\mathbf{W}[:, i] + \mathbf{s}$

**end for**

**for**  $i \in 1, \dots, L$  **do**

**if**  $i \leq \ell$  **then**

$m_i \leftarrow \text{MAX}(\hat{\mathbf{T}}[:, i])$

**else**

$m_i \leftarrow \text{MAX}(\hat{\mathbf{T}}[:, i]) + \mu$

**end if**

**end for**

$h^L \leftarrow \text{HISTOGRAM}(m_1, \dots, m_L)$

---

modeling assumption is reasonable and whether such a procedure will, in fact, recover the true statistics. Our two technical results answer these questions, and for brevity, we present them in the supplementary material accompanying the main paper. The informal summary

of the results is (1) the basic model (i.e., low-rank and low-variance residual) is indeed meaningful for the setting we are interested in, and (2) Recovering the low-rank and the residual by Algorithm 2 guarantees a high fidelity estimate of the max null, and shows that the error is small.

## 5. Experimental Setup

We evaluate RapidPT in multiple phases. First we perform a simulation study where the goal is to empirically demonstrate the requirements on the input hyperparameters that will guarantee an accurate recovery of the max null. These empirical results are compared to the analytical bounds governed by the theory (and further discussed in the supplement). The purpose of these evaluations is to walk the reader through the choices of hyperparameters, and how the algorithm is very robust to them. Next, we perform another simulation study where our goal is to evaluate the performance of RapidPT on multiple synthetic datasets generated by changing the strength of group-wise differences and the sparsity of the signal (e.g., how many voxels are different). We then conduct an extensive experiments to evaluate RapidPT against competitive methods on real brain imaging datasets. These include comparisons of RapidPT’s accuracy, runtime speedups and overall performance gains against two baselines. The first baseline we used was the latest release of the widely used MATLAB toolbox for nonparametric permutation testing in neuroimaging, Statistical NonParametric Mapping (SnPM) (SnPM (2013); Nichols and Holmes (2002)). The second baseline was a standard MATLAB implementation of algorithm 1, which we will call NaivePT. Both baselines serve to evaluate RapidPT’s accuracy. Further, the very small differences between the results provided by SnPM and NaivePT offer a secondary reference point that tells us an acceptable range for RapidPT’s results (in terms of differences). For runtime performance, SnPM acts as a state of the art baseline. On the other hand, NaivePT is used to evaluate how an unoptimized permutation testing implementation will perform on the datasets we use in our experiments. In the next section, we describe the experimental data, the hyperparameters space evaluated, the methods used to quantify accuracy, and the environment where all experiments were run.

### 5.1. Simulation Data I

The dataset consisted of  $n = 30$  synthetic images composed of  $v = 20000$  voxels. The signal in each voxel is derived from one of the following two normal distributions:  $N(\mu = 0, \sigma^2 = 1)$  and  $N(\mu = 1, \sigma^2 = 1)$ . Two groups were then constructed with 15 images in each and letting 1% (200 voxels) exhibit voxel-wise group differences. The signal in the remaining 99% of the voxels was assumed to come from  $N(\mu = 0, \sigma^2 = 1)$ .

### 5.2. Simulation Data II

The dataset consisted of a total of 48 synthetically generated datasets, each with  $v = 20000$  voxels. The datasets were generated by varying: the number of images ( $n$ ) in the dataset, the strength of the signal (i.e., deviation of  $\mu$  in  $N(\mu, 1)$  from  $N(0, 1)$ ) and the sparsity of the signal (percentage of voxels showing group differences). The dataset sizes were  $n = 60$ ,  $n = 150$ ,  $n = 600$ . Each dataset was split into two *equally sized* groups for which various degrees of signal and sparsity of the signal were used to generate the different



datasets. The first “group” (for group-difference analysis) in all datasets was generated from a standard normal distribution,  $N(\mu = 0, \sigma^2 = 1)$ . In the second “group”, we chose  $\{1\%, 5\%, 10\%, 25\%\}$  of the voxels from one of four normal distributions ( $N(\mu = 1, \sigma^2 = 1)$ ,  $N(\mu = 5, \sigma^2 = 1)$ ,  $N(\mu = 10, \sigma^2 = 1)$ ,  $N(\mu = 25, \sigma^2 = 1)$ ). The signal in the remaining voxels in the second group was also obtained from  $N(\mu = 0, \sigma^2 = 1)$ .

To summarize the simulation setup, Simulation Data I fixes the dataset and changes the algorithmic hyperparameters whereas Simulation Data II fixes the algorithmic hyperparameters and generates different datasets.

### 5.3. Data

The data used to evaluate RapidPT comes from the Alzheimer’s disease Neuroimaging Initiative-II (ADNI2) dataset. The ADNI project was launched in 2003 by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, the Food and Drug Administration, private pharmaceutical companies, and nonprofit organizations, as a \$60 million, 5-year public-private partnership. The overall goal of ADNI is to test whether serial MRI, positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of MCI and early AD. Determination of sensitive and specific markers of very early AD progression is intended to aid researchers and clinicians to develop new treatments and monitor their effectiveness, as well as lessen the time and cost of clinical trials. The principal investigator of this initiative is Michael W. Weiner, M.D., VA Medical Center and University of California — San Francisco. ADNI is the result of the efforts of many co-investigators from a broad range of academic institutions and private corporations, and subjects have been recruited from over 50 sites across the U.S. and Canada. The initial aim of ADNI was to recruit 800 adults, ages 55 to 90, to participate in the research — approximately 200 cognitively normal older individuals to be followed for 3 years, 400 people with MCI to be followed for 3 years, and 200 people with early AD to be followed for 2 years.

For the experiments presented in this paper, we used gray matter tissue probability maps derived from T1-weighted magnetic resonance imaging (MRI) data. From this data, we constructed four varying sized datasets. We sampled  $n_1$  and  $n_2$  subjects from the CN and AD groups in the cohort, respectively. Table 1 shows a summary of the datasets used for our evaluations.

Dataset Size: $n$ ( $n_1, n_2$ )			
50 (25,25)	100 (50,50)	200 (100,100)	400 (200,200)

Table 1: Dataset sizes used in our experiments. The table lists the total number of subjects ( $n$ ) and how many of the participants were sampled from the CN ( $n_1$ ) and AD groups ( $n_2$ ).

#### 5.3.1. Data Pre-processing

All images were pre-processed using voxel-based morphometry (VBM) toolbox in Statistical Parametric Mapping software (SPM, <http://www.fil.ion.ucl.ac.uk/spm>). After pre-processing, we obtain a data matrix  $\mathbf{X}$  composed of  $n$  rows and  $v$  columns for each dataset shown in Table 1. Each row in  $\mathbf{X}$  corresponds to a subject and each column is associated to a voxel that denotes approximately the same anatomical location across subjects (since the

images are already co-registered). This pre-processing is commonly used in the literature and not specialized to our experiments.

#### 5.4. Hyperparameters

As outlined in Algorithm 2, there are three high-level input parameters that will impact the performance of the procedure: the *number of training samples* ( $l$ ), the *sub-sampling rate* ( $\eta$ ), and the *number of permutations* ( $L$ ). To demonstrate the robustness of the algorithm to these parameter settings, we explored and report on hundreds of combinations of these hyperparameters on each dataset. This also helps us identify the general scenarios under which RapidPT will be a much superior alternative to regular permutation testing. The baselines for a given combination of these hyperparameters are given by the max null distribution constructed by SnPM and NaivePT. The number of permutations used for the max null distributions of the baselines is the same as the number of permutations used by RapidPT for a given combination of hyperparameters.

- *Number of Training Samples:* The number of training samples,  $l$ , determines how many columns of  $\mathbf{T}$  are calculated to estimate the basis of the subspace, and also how many training passes are performed to estimate the shift that corrects for the bias-variance tradeoff discussed in Section 4.1. We decided to use the total number of subjects  $n$  as a guide to pick a sensible  $l$ , the rationale is that the maximum possible rank of  $\mathbf{T}$  is  $n$  (as discussed in Section 3). Further,  $l$  is also used to determine the number of passes performed to calculate the shift of the max null distribution. Calculating the shift is a cheap step, therefore it makes sense to use all the information available in  $\mathbf{T}_{ex}$  to calculate the shift. Table 2 shows a summary of the different values for  $l$  used in our evaluations.

Number of Training Samples: $l$				
Simulations	$\frac{n}{3}$	$n$	$2n$	
Experiments	$\frac{n}{2}$	$\frac{3n}{4}$	$n$	$2n$

Table 2: Number of training samples used to evaluate RapidPT.  $n$  corresponds to the total number of subjects in the dataset. For instance, for the 400 subject dataset the values for  $l$  used were 100, 200, 400, and 800.

- *Sub-sampling rate:* The sub-sampling rate,  $\eta$ , is the percentage of all the entries of  $\mathbf{T}$  that we will calculate (i.e., sample) when recovering the max null distribution. In the recovery phase,  $\eta$  determines how many voxel-wise test statistics will be calculated at each permutation to recover a column of  $\bar{\mathbf{T}}$ . For instance, if the data matrix has  $v$  columns (number of voxels) then instead of calculating  $v$  test statistics, we will sample only  $\eta v$  (where  $\eta \ll 1$ ) random columns and calculate test statistics only for those columns. Table 3 shows a summary of the different values for  $\eta$  used in our evaluations.
- *Number of Permutations* The number of permutations,  $L$ , determines the total number of columns in  $\mathbf{T}$ . By varying  $L$  we are able to see how the size of  $\mathbf{T}$  affects the accuracy of the algorithm and also how it scales compared to a standard permutation testing

Sub-sampling rate: $\eta$									
<b>Simulations</b>	0.5%	1%	1.6%	2%	4%	8%	16%	32%	64%
<b>Experiments</b>	0.1%	0.35%	0.5%	0.7%	1.5%	5%			

Table 3: Sub-sampling rates used to evaluate RapidPT.  $\eta$  is the percentage of the total number of entries in  $\mathbf{T}$  that will be calculated during the recovery phase.

implementation with the same number of permutations (e.g., NaivePT or SnPM). Table 4 shows a summary of the different values for  $L$  used in our evaluations.

Number of Permutations: $L$							
<b>Simulations</b>	5,000	10,000	20,000	40,000	50,000	100,000	
<b>Experiments</b>	2,000	5,000	10,000	20,000	40,000	80,000	160,000

Table 4: Number of permutations done to evaluate RapidPT.

### 5.5. Accuracy Benchmarks

In order to assess the accuracy and overall usefulness of the recovered max null distribution by RapidPT we used three different measures: Kullback-Leibler Divergence (KL-Divergence),  $t$ -thresholds/ $p$ -values and the resampling risk.

*Kullback-Leibler Divergence:* The KL-Divergence provides a measure of the difference between two probability distributions. One of the distributions represents the ground truth (SnPM or NaivePT) and the other an “approximation” (obtained via RapidPT). In this case, the distributions are the max null distributions ( $h^L$ ). We use the KL-Divergence to identify under which circumstances (i.e., hyperparameters) RapidPT provides a good estimate of the overall max null distribution and if there are cases where the results are unsatisfactory.

*T-Thresholds/ $p$ -values:* Once we have evaluated whether all methods recover a similar max null distribution, we analyze if  $t$ -thresholds associated to a given  $p$ -value calculated from each max null distribution are also similar.

*Resampling Risk:* Two methods can recover a similar max null distribution and  $p$ -values, and yet partially disagree in *which voxels* should be classified as statistically significant (e.g., within a group difference analysis). The resampling risk is the probability that the decision of accepting/rejecting the null hypothesis differs between two methods (Jockel (1984)). Let  $v_1$  and  $v_2$  be the number of voxels whose null hypothesis was rejected according to the max null derived from Method 1 and 2, respectively. Further, let  $v_c$  be the number of voxels that are the (set) intersection of  $v_1$  and  $v_2$ . The resampling risk can then be calculated as shown in (6).

$$\text{risk} = \frac{\frac{v_1 - v_c}{v_1} + \frac{v_2 - v_c}{v_2}}{2} \quad (6)$$

### 5.6. Implementation Environment and other details

All evaluation runs reported in this paper were performed on multiple machines with the same hardware configuration. The setup consisted of multiple 16 core machine with two Intel(R) Xeon(R) CPU E5-2670, each with 8 cores. This means that any MATLAB application will be able to run a maximum of 16 threads. To evaluate the runtime performance of

RapidPT, we performed all experiments on two different setups which forced MATLAB to use a specific number of threads. First, we forced MATLAB to only use a single thread (single core) when running SnPM, RapidPT, and NaivePT. The performance results on a single threaded environment attempt to emulate a scenario where the application was running on an older laptop/workstation serially. In the second setup, we allow MATLAB to use all 16 threads available to demonstrate how RapidPT is also able to leverage a parallel computing environment to reduce its overall runtime.

Although all machines had the same hardware setup, to further ensure that we were making a fair runtime performance comparison, all measurements of a given Figure shown in Section 6 were obtained from the same machine.

## 6. Results

The hyperparameter space explored through hundreds of runs allowed identifying specific scenarios where RapidPT is most effective. To demonstrate accuracy, we first show the impact of the hyperparameters on the recovery of the max null distribution by analyzing KL-Divergence. Then we focus on the comparison of the corrected  $p$ -values across methods and the resampling risk associated with those  $p$ -values. To demonstrate the runtime performance gains of RapidPT, we first calculate the speedup results across hyperparameters. We then focus on the hyperparameters that a user would use and look at how RapidPT, SnPM, and NaivePT scale with respect to the dataset and the number of permutations. Overall, the large hyperparameter space that was explored in these experiments produced hundreds of figures. In this section, we summarize the results of all figures within each subsection, but only present the figures that we believe will convey the most important information about RapidPT. An extended results section is presented in the supplementary materials. We point out that following the results and the corresponding discussion of the plots and figures, we discuss the open-source toolbox version of RapidPT that is mde available online.

### 6.1. Accuracy

#### Simulations

Figure 4 shows the log KL-Divergence between the max null recovered by regular permutation testing and the max null recovered by RapidPT. We can observe that once the sub-sampling rate,  $\eta$ , exceeds the minimum value established by LRMC theory, RapidPT is able to accurately recover the max null with a KL-Divergence  $< 10^{-2}$ . Furthermore, increasing  $l$  can lead to slightly lower KL-Divergence as seen in the middle and right most plots. Finally, increasing the number of permutation improves the accuracy. A through discussion of how to choose the important hyperparameters is in Section 7.3.




Figure 4: KL-Divergence between the true max null and the one recovered by RapidPT. Each line corresponds to a different number of permutations. The dotted lines are the theoretical minimum sub-sampling rate and the "practical" one, i.e., the one the toolbox will set it to automatically if none is specified.

Figure 5 shows the log percent difference between the  $t$ -thresholds for different  $p$ -values obtained from the true max null and the one recovered by RapidPT. Similar to Figure 4, it is evident that once the strict requirement on the minimum value of  $\eta$  is achieved, we obtain a reliable  $t$ -threshold i.e., percent difference  $< 10^{-3}$ . Additionally, increasing the number of training samples (progression of plots from left to right) gives a improvement in the accuracy, however, not incredibly significant since we are already at negligible percent differences. Overall we see that RapidPT is able to estimate accurate thresholds even at extremely low  $p$ -value regimes.




Figure 5: Percent difference between the  $t$ -threshold (for different  $p$ -values) obtained from the true max null and the one recovered by RapidPT. The dotted lines are the theoretical minimum sub-sampling rate and the "practical" one, i.e., the one the toolbox will set it to automatically if none is specified.

Figure 6 shows the KL-Divergence between the max null recovered by regular permutation testing and the max null recovered by RapidPT on 48 synthetically generated datasets (16 in each column). The input hyperparameters used were fixed to  $L = 50000$ ,  $\eta = 2\eta_{min}$ , and  $l = n$ , where  $\eta_{min}$  refers to the theoretical minimum sub-sampling rate and  $n$  is the dataset size. As expected, the strength or sparsity of the signal does not have an impact on the performance of RapidPT. The dataset size, however, does have a slight impact on the accuracy but we still find that the recovered  $t$ -thresholds for the smaller datasets are within 2% of the true threshold.




Figure 6: KL-Divergence between the true max null and the one recovered by RapidPT on 48 datasets. The sub-sampling rate,  $\eta$ , used for each run was  $2\eta_{min}$ . The number of training samples,  $l$ , used for each run was  $n$  (i.e., the same as the number of images in the dataset).

## Results on the ADNI Dataset

*Can we recover the max null distribution?*

The left column of Figure 7 uses a colormap to summarize the KL-Divergence results obtained from comparing the max null distributions of a single run of SnPM versus multiple RapidPT runs with various hyperparameters. The right column of Figure 7 puts the numbers displayed in the colormaps into context by showing the actual max null distributions for a single combination of hyperparameters. Each row corresponds to each of the four datasets used in our evaluations.

The sub-sampling rate was the hyperparameter that had the most significant impact on the KL-Divergence. As shown in Figure 7, a sub-sampling rate of 0.1% led to high KL-Divergence, i.e., the max null distribution was not recovered in this case. For every other combination of hyperparameters RapidPT was able to sample at rates as low as 0.35% and still recover an accurate max null distribution. Most KL-Divergence values were in the 0.01 – 0.05 range with some occasional values between 0.05 – 0.15. However, using the max null distributions derived from only 2000 permutations leads to the resulting KL-Divergence to range mainly between 0.05 – 0.15, as shown in the supplementary results.

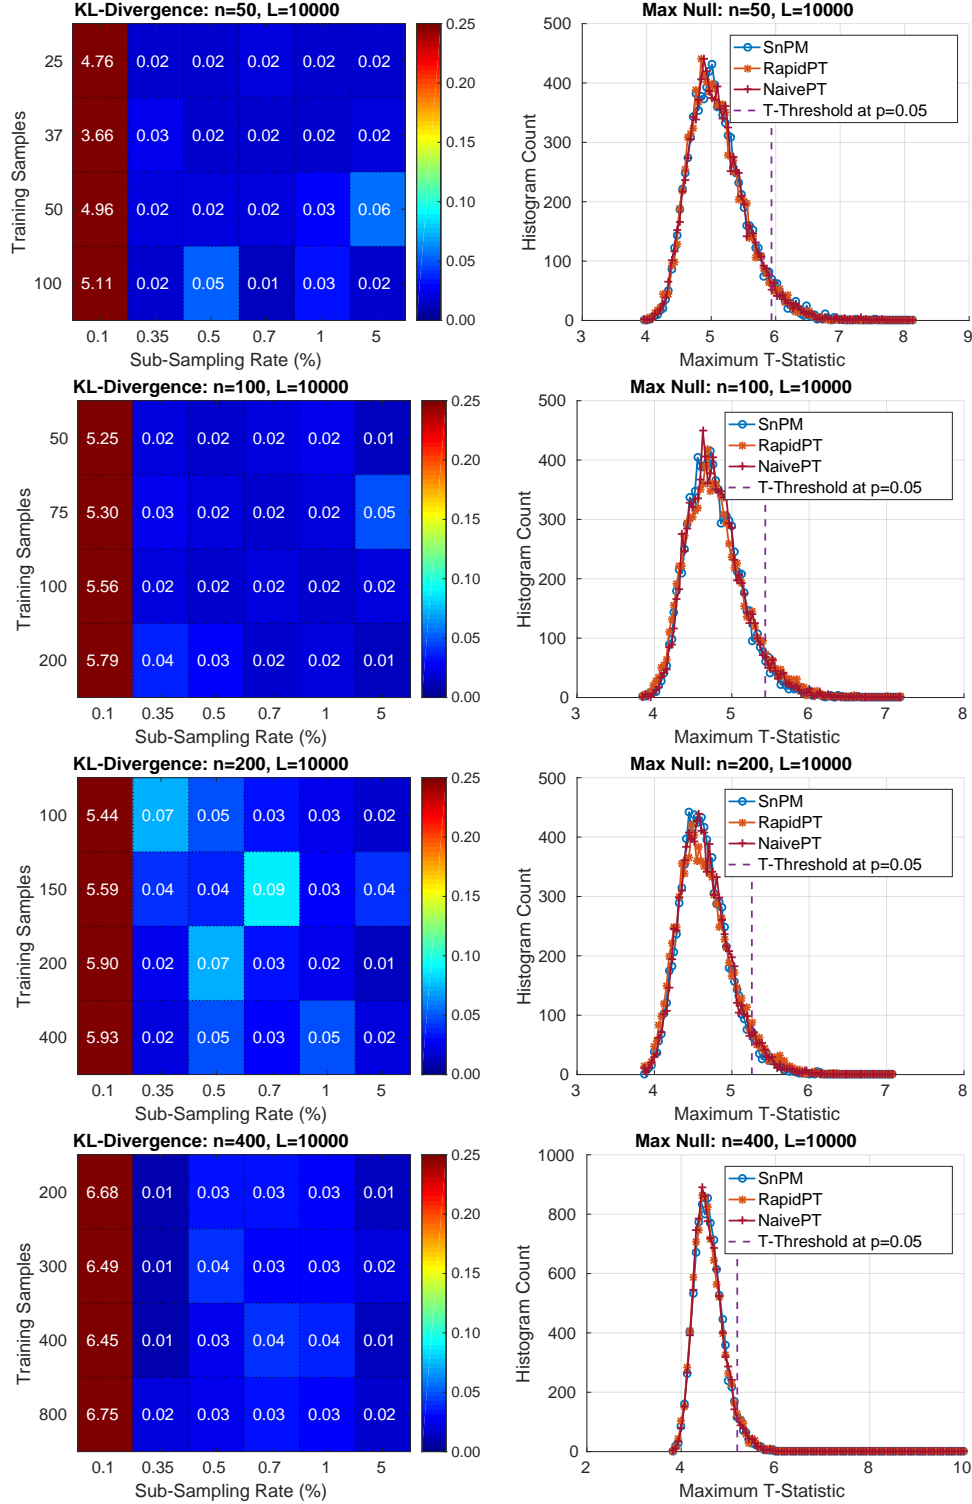


Figure 7: *Left*: Colormap of the KL-Divergence between the max null distributions of RapidPT and SnPM. Each colormap is associated to a run on one of the datasets and a fixed number of permutations. The resulting KL-Divergence from 24 hyperparameter combinations is displayed on each colormap. Rows 1, 2, 3, and 4 of this figure are associated to the 50, 100, 200, and 400 subject datasets respectively. *Right*: The max null distributions given by RapidPT, SnPM, and NaivePT for the hyperparameters:  $L = 10000$ ,  $l = n$ , and  $\eta = 0.35\%$ .

*Are we rejecting the correct null hypotheses?*

The test statistics obtained using the original data labels whose value exceed the  $t$ -threshold associated to a given  $p$ -value will correspond to the null hypothesis rejected. Figure 8 shows the resultant mapping between  $t$ -threshold and  $p$ -values for the max null distribution for a given set of hyperparameters. It is evident that the difference across methods is minimal. Moreover, Figure 8 shows that low  $p$ -values ( $p < 0.1$ ), which are the main object of interest, show the lowest differences. However, despite the low percent differences between the  $p$ -values, in the larger datasets (100, 200, and 400 subjects) RapidPT consistently yields slightly more conservative  $p$ -values near the tails of the distribution. Nonetheless, Figure 11 shows that the resampling risk between RapidPT and the two baselines remains very close to the resampling risk between both baselines. In practice, these plots show that RapidPT will reject the null hypothesis for a slightly lower number of voxels than SnPM or NaivePT.

Despite the slight difference in thresholds, the actual brain regions whose null hypotheses were rejected consistently match between both methods as shown in Figures 9 and 10. Additionally, the regions picked up by both RapidPT and SnPM in Figure 9 correspond to the *Hippocampus* – which is one of the primary structural brain imaging region that corresponds to the signature of cognitive decay at the onset of Alzheimer’s disease. The regions in Figure 10 contain a subset of the brain regions in Figure 9 which is expected from the thresholds shown in the right column of Figure 7.

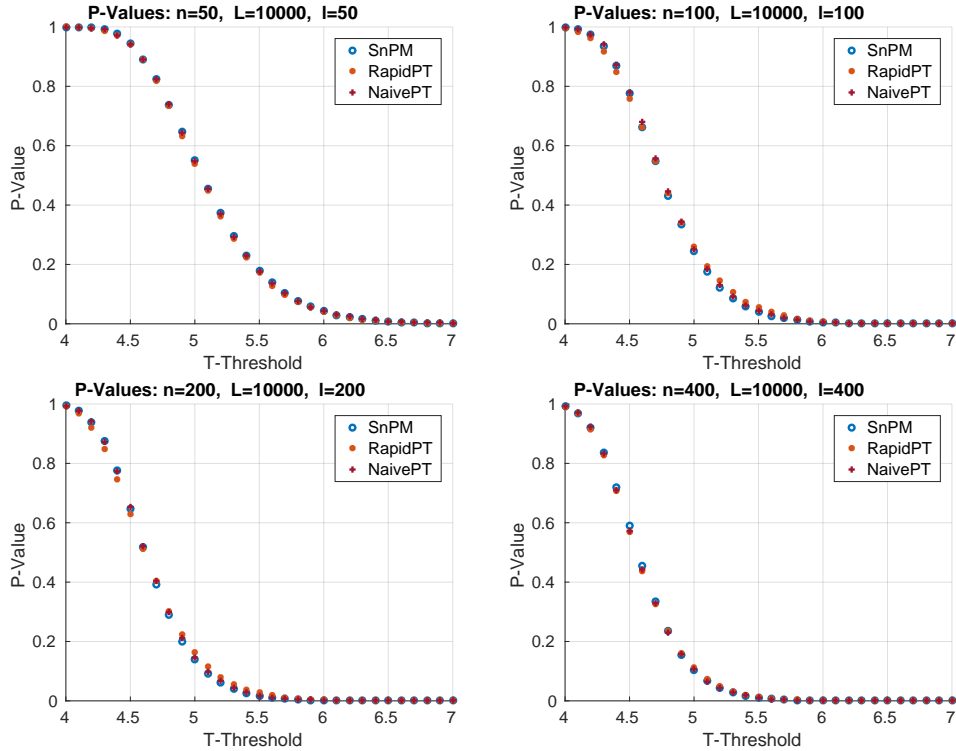




Figure 8:  $p$ -values for SnPM, RapidPT, and NaivePT. The hyperparameters used were:  $\eta = 0.35\%$ ,  $L = 10000$ , and  $l = n$ . The results in this plot were obtained from the max null distributions shown in the right hand side of Figure 7.



## ADNI Statistic Maps



(a) SnPM Statistic Maps. Slice -44 to 58 from left to right top to bottom.



(b) RapidPT Statistic Maps. Slice -44 to 58 from left to right top to bottom.  $\eta = 0.5\%$ ,  $l = 200$ .

cluster-level		voxel-level				x,y,z mm		
	$k$	$p_{FWE-corr}$	$p_{FDR-corr}$	$T$	$p_{uncorr}$			
<b>82</b>	<b>0.0000</b>	<b>0.0005</b>	<b>10.03</b>	0.0000		<b>-26</b>	<b>-12</b>	<b>-32</b>
	0.0000	0.0005	9.58	0.0000		-24	-24	-21
	0.0015	0.0005	5.84	0.0000		-30	-18	-26
<b>68</b>	<b>0.0000</b>	<b>0.0005</b>	<b>9.55</b>	0.0000		<b>27</b>	<b>-22</b>	<b>-21</b>
	0.0000	0.0005	8.96	0.0000		26	-10	-32
	<b>1372</b>	<b>0.0000</b>	<b>0.0005</b>	<b>9.04</b>	0.0000	<b>-33</b>	<b>-9</b>	<b>-20</b>
<b>173</b>	0.0000	0.0005	8.38	0.0000		-38	-28	-9
	0.0000	0.0005	7.97	0.0000		12	-6	27
	<b>173</b>	<b>0.0000</b>	<b>0.0005</b>	<b>8.74</b>	0.0000	<b>36</b>	<b>-6</b>	<b>-22</b>
<b>21</b>	0.0000	0.0005	7.85	0.0000		38	-34	-4
	0.0000	0.0005	7.51	0.0000		39	-15	-16
	<b>75</b>	<b>0.0000</b>	<b>0.0005</b>	<b>7.67</b>	0.0000	<b>-21</b>	<b>-36</b>	<b>3</b>
<b>15</b>	<b>0.0000</b>	<b>0.0005</b>	<b>7.59</b>	0.0000		<b>14</b>	<b>9</b>	<b>0</b>
	0.0000	0.0005	6.55	0.0000		9	4	-6
	0.0001	0.0005	6.25	0.0000		18	12	6
<b>23</b>	<b>0.0000</b>	<b>0.0005</b>	<b>7.51</b>	0.0000		<b>0</b>	<b>0</b>	<b>2</b>
	<b>0.0000</b>	<b>0.0005</b>	<b>6.76</b>	0.0000		<b>28</b>	<b>-40</b>	<b>-6</b>
	0.0000	0.0005	6.55	0.0000		26	-46	0

Height threshold: statistic  $u = 5.16$  (0.0500 FWE) Degrees of freedom = [1 398]  
 Design: 2 Groups: Two Sample T test; 1 scan per subject: 200(GrpA),200(GrpB)  
 Search vol: 1939410 cmm, 574640 voxels  
 Perms: 100000 permutations of conditions, bhPerms=0 Voxel size: [ 1.50, 1.50, 1.50] mm

(c) SnPM Report


cluster-level		voxel-level				x,y,z mm		
	$k$	$p_{FWE-corr}$	$p_{FDR-corr}$	$T$	$p_{uncorr}$			
<b>79</b>	<b>0.0000</b>	<b>0.0004</b>	<b>10.03</b>	0.0000		<b>-26</b>	<b>-12</b>	<b>-32</b>
	0.0000	0.0004	9.58	0.0000		-24	-24	-21
	0.0042	0.0004	5.84	0.0000		-30	-18	-26
<b>65</b>	<b>0.0000</b>	<b>0.0004</b>	<b>9.55</b>	0.0000		<b>27</b>	<b>-22</b>	<b>-21</b>
	0.0000	0.0004	8.96	0.0000		26	-10	-32
	<b>1218</b>	<b>0.0000</b>	<b>0.0004</b>	<b>9.04</b>	0.0000	<b>-33</b>	<b>-9</b>	<b>-20</b>
<b>164</b>	0.0000	0.0004	8.38	0.0000		-38	-28	-9
	0.0000	0.0004	7.97	0.0000		12	-6	27
	<b>164</b>	<b>0.0000</b>	<b>0.0004</b>	<b>8.74</b>	0.0000	<b>36</b>	<b>-6</b>	<b>-22</b>
<b>20</b>	0.0000	0.0004	7.85	0.0000		38	-34	-4
	0.0000	0.0004	7.51	0.0000		39	-15	-16
	<b>63</b>	<b>0.0000</b>	<b>0.0004</b>	<b>7.67</b>	0.0000	<b>-21</b>	<b>-36</b>	<b>3</b>
<b>15</b>	<b>0.0000</b>	<b>0.0004</b>	<b>7.59</b>	0.0000		<b>14</b>	<b>9</b>	<b>0</b>
	0.0001	0.0004	6.55	0.0000		9	4	-6
	0.0006	0.0004	6.25	0.0000		18	12	6
<b>9</b>	<b>0.0000</b>	<b>0.0004</b>	<b>7.51</b>	0.0000		<b>0</b>	<b>0</b>	<b>2</b>
	<b>0.0000</b>	<b>0.0004</b>	<b>7.45</b>	0.0000		<b>-27</b>	<b>-40</b>	<b>-8</b>
	<b>21</b>	<b>0.0001</b>	<b>0.0004</b>	<b>6.76</b>	0.0000	<b>28</b>	<b>-40</b>	<b>-6</b>
	0.0001	0.0004	6.55	0.0000		26	-46	0

Height threshold: statistic  $u = 5.25$  (0.0500 FWE) Degrees of freedom = [1 398]  
 Design: 2 Groups: Two Sample T test; 1 scan per subject: 200(GrpA),200(GrpB)  
 Search vol: 1939410 cmm, 574640 voxels  
 Perms: 100000 permutations of conditions, bhPerms=0 Voxel size: [ 1.50, 1.50, 1.50] mm


(d) RapidPT Report

Figure 9: Thresholded FWER corrected statistical maps at ( $\alpha = 0.05$ ) with the  $n = 400$  dataset. The hyperparameters used were:  $\eta = 0.5\%$ ,  $l = n$ , and  $L = 100000$ . The images show the test statistics for which the null was rejected in SnPM (top) and RapidPT (bottom). The tables show a numerical summary of the images. The columns refer to:  $k$  - cluster size,  $p_{FWE-corr}$  - corrected  $p$ -values,  $T$  - max cluster t-statistic.  $p_{FWE-corr}$  that appears as 0.0000 are  $1e^{-5}$ . These results were obtained from a run with the hyperparameters specified in the title.

## ADNI Statistic Maps



(a) SnPM P-Map. Slice -44 to 58 from left to right top to bottom.



(b) RapidPT P-Map. Slice -44 to 58 from left to right top to bottom.  $\eta = 0.5\%$ ,  $l = 200$ .

cluster-level		voxel-level				x,y,z mm		
	<i>k</i>	$p_{FWE-corr}$	$p_{FDR-corr}$	<i>T</i>	$p_{uncorr}$			
	6	0.0000	0.0111	7.07	0.0000	36	-6	-22
	24	0.0001	0.0111	6.49	0.0000	-24	-24	-21
		0.0004	0.0111	6.29	0.0000	-26	-16	-28
		0.0028	0.0111	5.89	0.0000	-24	-9	-33
	2	0.0001	0.0111	6.48	0.0000	26	-10	-32
	9	0.0001	0.0111	6.44	0.0000	27	-24	-20
	1	0.0008	0.0111	6.14	0.0000	-9	3	24
	1	0.0023	0.0111	5.93	0.0000	-38	-28	-9
	2	0.0028	0.0111	5.89	0.0000	-20	-45	4
	2	0.0030	0.0111	5.87	0.0000	2	18	10
	1	0.0037	0.0111	5.83	0.0000	-8	0	24
	2	0.0044	0.0111	5.79	0.0000	2	15	14
	1	0.0051	0.0111	5.75	0.0000	14	-15	28
	1	0.0054	0.0111	5.74	0.0000	12	-6	27
	1	0.0066	0.0111	5.70	0.0000	-8	8	21
	1	0.0074	0.0111	5.68	0.0000	6	-12	-16
	1	0.0116	0.0111	5.58	0.0000	40	-15	-18
	2	0.0157	0.0111	5.52	0.0000	26	-8	-33

Height threshold: statistic  $u = 5.25$  (0.0500 FWE) Degrees of freedom = [1 198]  
 Design: 2 Groups: Two Sample T test; 1 scan per subject: 100(GrpA),100(GrpB)  
 Search vol: 1.917570e+06 cmm, 568169 voxels  
 Perms: 100000 permutations of conditions, bhPerms=0 Voxel size: [ 1.50, 1.50, 1.50] mm

(c) SnPM Report

cluster-level		voxel-level				x,y,z mm		
	<i>k</i>	$p_{FWE-corr}$	$p_{FDR-corr}$	<i>T</i>	$p_{uncorr}$			
	6	0.0000	0.0008	7.07	0.0000	36	-6	-22
	21	0.0002	0.0008	6.49	0.0000	-24	-24	-21
		0.0007	0.0008	6.29	0.0000	-26	-16	-28
		0.0042	0.0008	5.89	0.0000	-24	-9	-33
	2	0.0002	0.0008	6.48	0.0000	26	-10	-32
	9	0.0003	0.0008	6.44	0.0000	27	-24	-20
	1	0.0014	0.0008	6.14	0.0000	-9	3	24
	1	0.0036	0.0008	5.93	0.0000	-38	-28	-9
	2	0.0044	0.0008	5.89	0.0000	-20	-45	4
	2	0.0046	0.0008	5.87	0.0000	2	18	10
	1	0.0057	0.0008	5.83	0.0000	-8	0	24
	2	0.0067	0.0008	5.79	0.0000	2	15	14
	1	0.0080	0.0008	5.75	0.0000	14	-15	28
	1	0.0082	0.0008	5.74	0.0000	12	-6	27
	1	0.0099	0.0008	5.70	0.0000	-8	8	21
	1	0.0108	0.0008	5.68	0.0000	6	-12	-16
	1	0.0164	0.0008	5.58	0.0000	40	-15	-18
	2	0.0213	0.0008	5.52	0.0000	26	-8	-33

Height threshold: statistic  $u = 5.30$  (0.0500 FWE) Degrees of freedom = [1 198]  
 Design: 2 Groups: Two Sample T test; 1 scan per subject: 100(GrpA),100(GrpB)  
 Search vol: 1.917570e+06 cmm, 568169 voxels  
 Perms: 100000 permutations of conditions, bhPerms=0 Voxel size: [ 1.50, 1.50, 1.50] mm

(d) RapidPT Report

Figure 10: Thresholded FWER corrected statistical maps at ( $\alpha = 0.05$ ) with the  $n = 200$  dataset. The hyperparameters used were:  $\eta = 0.5\%$ ,  $l = n$ , and  $L = 100000$ . The images show the test statistics for which the null was rejected in SnPM (top) and RapidPT (bottom). The tables show a numerical summary of the images. The columns refer to:  $k$  - cluster size,  $p_{FWE-corr}$  - corrected  $p$ -values,  $T$  - max cluster t-statistic.  $p_{FWE-corr}$  that appears as 0.0000 are  $1e^{-5}$ . These results were obtained from a run with the hyperparameters specified in the title.

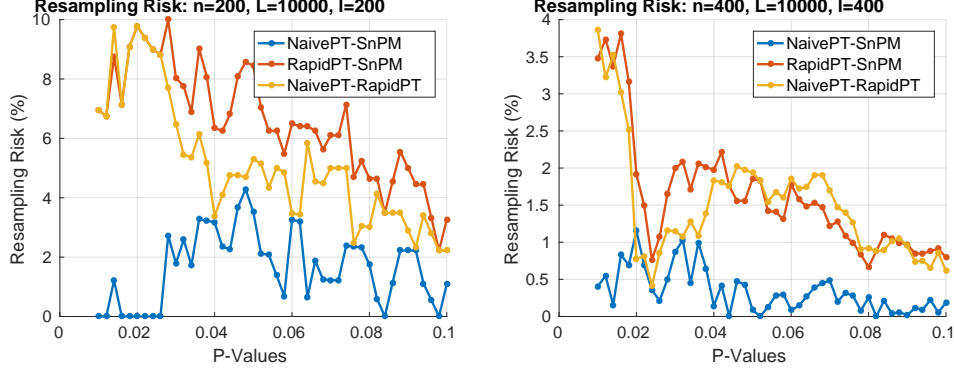


Figure 11: Resampling risk of NaivePT-SnPM, RapidPT-SnPM, and NaivePT-RapidPT. The hyperparameters used were:  $\eta = 0.35\%$ ,  $L = 10000$ , and  $l = n$ .

## 6.2. Runtime Performance

### 6.2.1. Effect of hyperparameters on the speed of RapidPT

Figures 12 and 13 show the speedup gains of RapidPT over SnPM and NaivePT, respectively. Each column corresponds to a single dataset, and each row corresponds to a different number of permutations. The supplementary results include an exhaustive version of these results that show the speedup gains of RapidPT for many additional number of permutations.

As shown in Figure 12, RapidPT outperforms SnPM in most scenarios. With the exception of the  $L = 2000$  and  $L = 5000$  runs on the larger datasets ( $n = 200$  and  $n = 400$ ), the colormaps show that RapidPT is 1.5-30x faster than SnPM. As expected, a low  $\eta$  (0.35%, 0.5%) and  $l$  ( $\frac{n}{2}, \frac{3n}{4}$ ) leads to the best runtime performance without a noticeable accuracy tradeoff, as can be seen also in Fig. 7 earlier.

Figure 13 shows how RapidPT performs against a non-optimized permutation testing implementation. In this setup, RapidPT outperforms NaivePT in every single combination of the hyperparameters. The same speedup trends that were seen when comparing RapidPT and SnPM are seen between RapidPT and NaivePT but with a much larger magnitude. For the remainder of this section, the runtime results of NaivePT are no longer compared because the difference is too large.




Figure 12: Colormaps of the speedup gains of RapidPT over SnPM in a serial and parallel computing environment. Each colormap corresponds to a run with a given dataset and a fixed number of permutations, and displays 20 different speedups resulting from different hyperparameter combinations. The first two rows correspond to the speedups obtained from the runs on 16 cores, and the last two columns from runs on a single core. Columns 1, 2, 3, and 4 of this figure correspond to the 50, 100, 200, and 400 subject datasets, respectively.

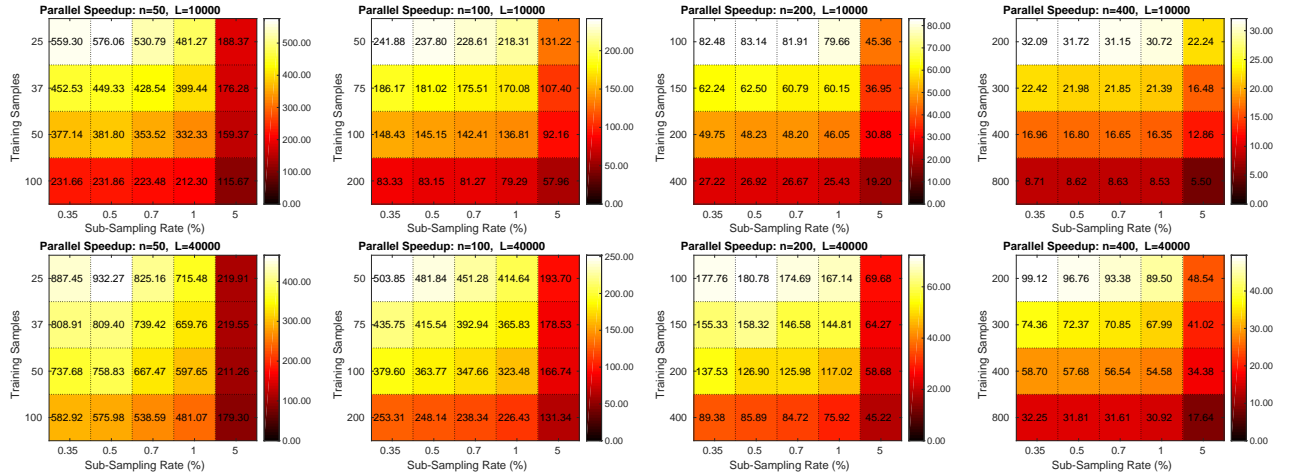


Figure 13: Colormap of the speedup gains of RapidPT over NaivePT. The organization of the colormaps and the information they display is the same as Figure 12. These speedups correspond to both programs running on a MATLAB instance that could use all 16 available cores.



### 6.2.2. Scaling of RapidPT vs. SnPM

As opposed to  $\eta$  and  $l$  which only seem to have an impact on the runtime performance of RapidPT, the number of permutations and the size of the dataset have an impact on the runtime of both RapidPT and SnPM. In this section, we compare how RapidPT and SnPM scale as we vary these two parameters. Figures 14 and 15 show the runtime performance of RapidPT for  $\eta = 0.35\%$  and  $l = n$ .

*Number of Permutations:* Figure 14 shows the super linear scaling of SnPM compared to RapidPT for all datasets as the number of permutation increases. Doubling the number of permutations in SnPM leads to an increase in the runtime by a factor of about two. On the other hand, doubling the number of permutations for RapidPT only affects the runtime of the recovery phase leading to small increases in the timing performance. The performance of both implementations is only comparable if we focus on the lower range of the number of permutations ( 5000) across datasets. As this number increases, as was shown in Figure 12, RapidPT outperforms the permutation testing implementation within SnPM.

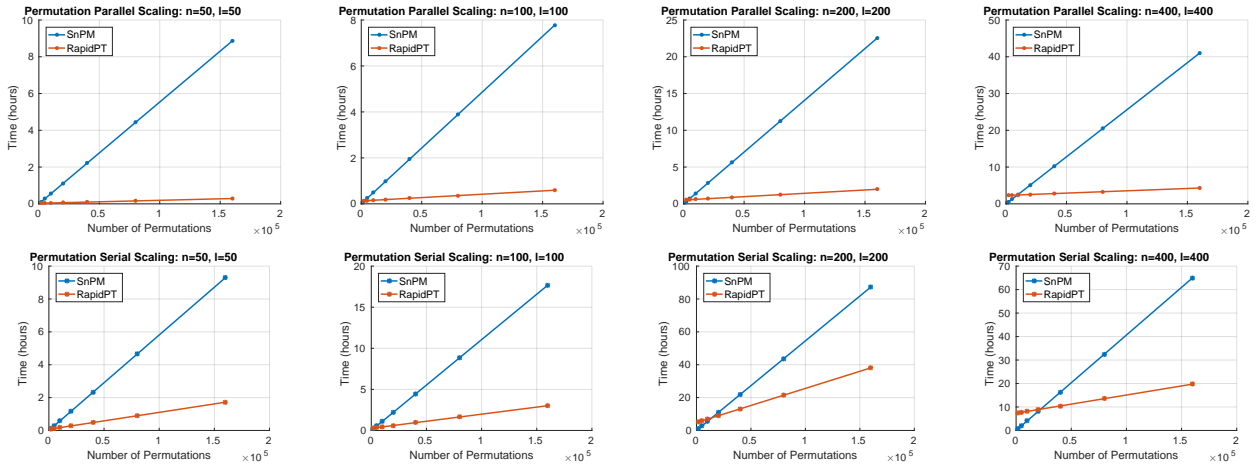


Figure 14: Effect of the number of permutations on the runtime performance of RapidPT and SnPM on 16 cores (first row) and on a single core (second row). The hyperparameters used were:  $\eta = 0.35\%$  and  $l = n$ .

*Dataset Size:* Figure 15 shows the effect of the dataset size on the runtime of RapidPT and SnPM. In SnPM, as expected, increasing the dataset by a factor of two leads to an increase on the runtime by a factor of two. In RapidPT, however, increasing the dataset size has a variable effect on the runtime. The training phase ends up contributing more to the runtime as the dataset size increases, while the recovery phase runtime increases at much slower rate.

Overall, scaling the number of permutations have a stronger impact on the runtime performance of SnPM than RapidPT. On the other hand, scaling the dataset size has a more negative effect on the timing of RapidPT than in SnPM. Furthermore, if both parameters are increased at the same time the runtime of SnPM increases at a much faster rate than the runtime of RapidPT.




Figure 15: Effect of the dataset size on the performance of RapidPT and SnPM on 16 cores (first row) and on a single core (second row). The overall measured time of RapidPT is the result of the total time spent on the training phase and the recovery phase. The hyperparameters used were:  $\eta = 0.35\%$ ,  $L = 10000$ , and  $l = n$ .

## 7. Discussion

### 7.1. Accuracy

#### 7.1.1. Recovered Max null Distribution

Monte carlo permutation tests perform a randomization step where a random subset of the total number of permutations is chosen. This means that the constructed max null distribution from one run might slightly differ from another run, and as the number of permutations increases, this difference will decrease. In terms of KL-Divergence, this means that the KL-Divergence between two permutation testing runs on the same data will be small, but *not exactly zero*. The results show that given a good set of hyperparameters the KL-Divergence between a run of RapidPT versus a regular permutation testing run leads to a very low KL-Divergence which is the expected result, even if we run the *same permutation testing program twice*. The evaluated scenarios show that a sensible set of hyperparameters can be easily defined as long as the sub-sampling rate is sufficient for the recovery of the permutation testing matrix. The minimum number of sub-sampled entries needed to accurately recover  $\mathbf{T}$  depends on the rank and dimensions of  $\mathbf{T}$  as discussed in section 3.1 and in the supplement. For the simulation study with  $n = 30$ , the minimum  $\eta$  required was 1.6% of all entries as shown in Figures 4 and 5. The experiments on the other 48 synthetic datasets (Figure 6) used  $2\eta_{min}$  as the sub-sampling rate. In our experiments on the ADNI dataset,  $\eta \geq 0.35\%$  led to a large enough set of sub-sampled entries to obtain an

accurate estimate of the max null distribution. For a brief discussion on how to pick a sensible  $\eta$  and the minimum sub-sampling rate ( $\eta_{min}$ ), please refer to section 3 of the supplementary material. Overall, once we have a sensible  $\eta$ , the resulting max null distribution constructed by RapidPT is consistent to the one recovered by regular permutation testing.

The number of permutations also has a significant impact on the KL-Divergence. As  $L$  increases, the KL-Divergence decreases. However, this is also true between any two permutation testing runs on the same dataset.

### 7.1.2. Evaluating $p$ -values

As seen in the  $p$ -value spectrum plots, the  $p$ -values drawn from each max null distribution agree (given a good set of hyperparameters as discussed above). This follows from the low KL-Divergence, since the KL-Divergence is a measure of the overall difference between the distributions. The lowest differences, however, are located in the tails of the distributions. This means that the derived thresholds are expected to accept/reject almost the same null hypotheses.

*The extremely low  $p$ -value regime:* As we will discuss in section 7.2 the largest speedups from RapidPT are obtained as the number of permutations increases. The main reason to increase  $L$  is to obtain smaller  $p$ -values. Figures 5 and 8 show that RapidPT recovers extremely accurate t-threshold (percent differences  $< 0.1\%$ ). Figures 9 and 10 demonstrate this in a practical scenario where a large number of permutation ( $L = 100,000$ ) and a high percentage of the brain regions that rejected the null had a  $p$ -value  $< 0.001$ . Therefore, a user interested in using small  $p$ -values can benefit from large computational speedups when using RapidPT, with a negligible loss in accuracy.

### 7.1.3. Resampling Risk

*Small signal datasets:* Although the  $p$ -values agree across methods, a small difference may lead to slightly more or less null hypotheses to be rejected by a given method. This has a direct impact on the resampling risk between RapidPT and regular permutation testing. In datasets where there is a small signal difference between groups, we may see an elevated resampling risk. The reason is because a very small number of null hypotheses will be rejected. Therefore, at a given  $p$ -value,  $p$ , one of the methods will reject a small number of null hypotheses which the second method will not reject until the  $p$ -value is  $p + \delta$ , where  $\delta \ll 1$ . This slight difference in the number of rejected null hypotheses may have a significant impact in the resampling risk. For instance in Figure 11 ( $N = 200$  at  $p = 0.05$ ), RapidPT rejected 59 (out of  $\sim 568k$  statistics) null hypotheses and SnPM rejected 71: this led to a resampling risk of 8.45%. On the other hand, for  $N = 400$  at  $p = 0.05$  RapidPT rejected 2158 (out of  $\sim 570k$  statistics) and SnPM rejected 2241 null hypotheses, resulting in a lower resampling risk of 1.85% even when there is a larger difference in the number of rejected hypotheses. Nonetheless, as we can see in the brain maps in Figures 9 and 10 that this difference of RapidPT's and SnPM's rejections were among the boundary voxels of the rejection region (i.e., the mismatch is not at the center of the rejection/significant region). Note that once a reasonable small smoothing filter is applied to nullify noise (e.g., stand-alone voxels) this apparent small difference will vanish visually. A similar situation is encountered in the 50 and 100 subject dataset (see supplement), where the number of null hypotheses rejected was extremely low ( $< 10$ ) and single mismatch led to an elevated resampling risk. Therefore,

the resampling risk is a useful measure if it is presented together with the number of null hypothesis being rejected. Hence, RapidPT yields a slightly more conservative test, primarily because it is based on sub-sampling. Nevertheless, this sub-sampling is robust to locating the same voxel clusters that displayed group differences as the other methods.

## 7.2. Runtime Performance

Our results show that under certain scenarios RapidPT provides substantial speedups over SnPM (state of the art) and NaivePT (simple implementation). From the runtime performance of NaivePT, it is evident that in practice it is more beneficial for the user to rely on a permutation testing toolbox such as SnPM. In the remainder of the discussion we will just consider the runtime performance of RapidPT and SnPM.

*Dataset Size:* Overall, the largest speedups in both the serial and parallel setups were obtained in the runs for the smallest dataset ( $N = 50$ ). The number of training samples used to estimate the basis  $\mathbf{U}$  is smaller and consequently the training phase time decreases. As the dataset size increases, the training time introduces a considerable overhead which negatively impacts the speedup of RapidPT over SnPM when less than 20000 permutations are being performed.

*Number of Permutations:* The number of permutations have a linear impact on the runtime of both RapidPT and SnPM. But as shown in Figure 14, RapidPT runtime increases at a lower rate than SnPM’s. This is expected since the number of permutations only impacts the runtime of the recovery phase. Therefore, the training phase time for a given setup is constant as the number of permutations changes. The results show that for datasets with less than or equal to 400 subjects between 5000 – 10000 permutations is the threshold where despite the training overhead in RapidPT, the expected speedup is considerable to justify its use. In the larger datasets, when performing less than 5000 permutations, the training overhead becomes too large, as shown in the supplementary material.

### 7.2.1. Serial vs. Parallel Performance

The serial and parallel runs show very similar speedup trends across hyperparameters as shown in the results and supplementary material. However, in terms of actual runtime, both RapidPT and SnPM benefit from being able to run on a parallel environment. This is an essential feature for any software toolbox that will be running on modern workstations because multiple cores are available in nearly all computers shipped today.

## 7.3. Hyperparameter Recommendations

*Sub-sampling Rate:* The KL-Divergence results, Figures 4 and 7, show that as long as the sub-sampling rate is greater than or equal to a certain threshold, RapidPT is able to accurately recover the max null distribution. For the simulation study this threshold was 1.6% and for the real data experiments 0.35% was a high enough sub-sampling rate. A minimum  $\eta_{min}$  can be calculated using LRMC theory as shown in Section 3 of the supplementary material. This minimum value is simply a function of the number of voxels and the number of data instances. The simulation results in Figures 4 and 5 explicitly show this  $\eta_{min}$ . The toolbox itself sets it to a default conservative value of  $2 * \eta_{min}$ , which is also shown in Figures 4 and 5. This slightly larger choice ensures the error in recovering the null is almost negligible. Overall, the accuracy of RapidPT will not significantly improve



as the sub-sampling rate increases (pass  $\eta_{min}$ ). On the other hand, a low sub-sampling rate can significantly reduce the runtime, in particular in the large  $L$  regime, and is therefore preferable. A user does not need to change this rate in practice (beyond what is given by the toolbox). However, if he/she is willing to be even more flexible by sacrificing some accuracy to achieve an even higher speed-up, it can be reduced appropriately.

*Number of Training Samples:* The number of training samples will ideally be the exact rank of  $\mathbf{T}$ . This is usually not known, however, we know that the rank is bounded by the number of subjects in the data matrix used to generate  $\mathbf{T}$ . Therefore, in our evaluations, we are able to accurately recover the max null distribution even when using as low as  $\frac{n}{2}$  training samples. The recommended number of training samples in practice is  $n$  (which follows from rank structure of the testing matrix) and is the default setting within RapidPT toolbox.

*Number of Permutations:* When a large number of permutations is desired, RapidPT should be strongly considered due to its runtime gains. Not only RapidPT provides considerable runtime gains in the large  $L$  regime, but also its accuracy will improve as the KL-Divergence results show. Note, however, that the user should also take into account the dataset size and look at the speedup colormaps to see the expected speedups.

*Dataset size:* Although not a hyperparameter, as discussed above, the dataset size should play a role when the user selects which method to use. Large datasets ( $n \geq 200$ ), RapidPT will be a good option if the user is planning to perform a large number of permutations. For medium-sized datasets ( $50 \leq n \leq 200$ ), RapidPT will most likely lead to good speedup gains. For small datasets ( $n \leq 50$ ), although RapidPT might lead to speedup gains over regular permutation testing, the total runtime will be on the order of minutes anyway.

As we briefly discussed in Section 1, Winkler et al. (2016) provides several strategies for reducing the runtime of permutation testing. But the authors do not report significant speedup gains against regular permutation testing on a 50 subject dataset with  $\approx 200000$  voxels. On the other hand, RapidPT is able to consistently outperform a state of the art permutation testing implementation (SnPM) on a 50 subject dataset with  $\approx 540k$  voxels. This boost in performance is, in large part, due to our low sub-sampling rates. Our subspace tracking algorithm is, nonetheless, able to perform recovery in this sparse sampling setting.

#### 7.4. SnPM Integration

SnPM is a toolbox that can be used within the software Statistical Parametric Mapping (SPM) SPM (2012). RapidPT has been integrated into the development version of SnPM SnPM (2013). This enables users to leverage the pre and post processing capabilities of SnPM. Through the graphical user interface (GUI) of SnPM the user can simply specify if they want to use RapidPT or not. Alternatively, the experienced user can also toggle a flag called *RapidPT* inside the *snpm\_defaults.m*. Once this flag is set the user can simply proceed with their normal SnPM workflow. The SnPM GUI does not allow the user to set RapidPT's hyperparameters ( $\eta$ ,  $l$ ), however, the online documentation walks the user through the process of setting them manually. A preview of the online documentation can be seen in Figure 16. Further discussion and walkthroughs of how to use SnPM and RapidPT within SnPM can be found in the documentation of both toolboxes Gutierrez-Barragan and Ithapu (2016); SnPM (2013). The RapidPT library webpage is at <http://felipegb94.github.io/RapidPT/>.

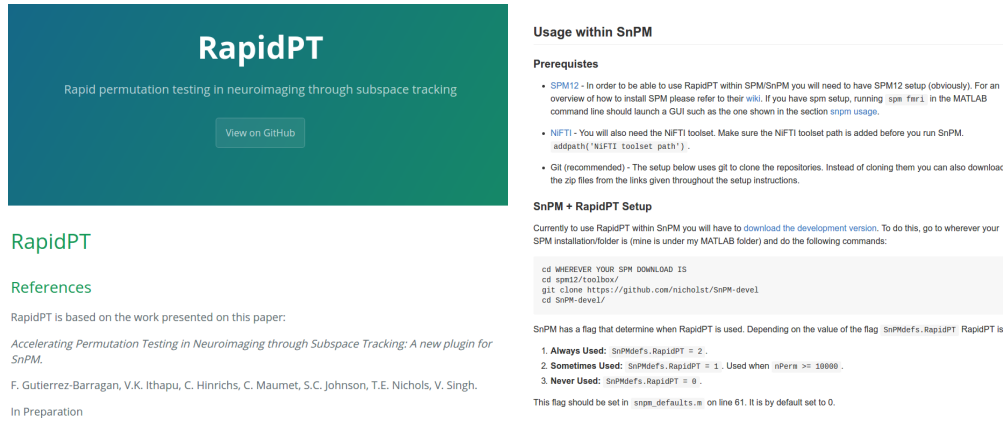


Figure 16: Screenshots of the RapidPT website (left) and SnPM integration documentation (right).

## 8. Conclusion

In this paper, we have presented a new algorithmic framework that is able to efficiently approximate the max null distribution commonly obtained through permutation testing. By exploiting the structure of the permutation testing matrix,  $\mathbf{T}$ , and applying recent ideas from online matrix completion we show through our theoretical analysis and experimental evaluations that one can subsample entries of  $\mathbf{T}$  at extremely low-rates and still construct a good estimate of  $\mathbf{T}$ . The algorithm first goes through a training phase where the basis and the distribution of the residual of  $\mathbf{T}$  are estimated. Then it continues into the recovery phase where a small number of entries of each column of  $\mathbf{T}$  are calculated and the rest are estimated through matrix completion. Finally, we obtain the max null distribution from the maximum value of each column in  $\mathbf{T}$ . Experiments on four varying sized datasets derived from the ADNI2 dataset showed that if we sub-sample at a high enough rate we can accurately recover the max null distribution at a fraction of the time in many scenarios. The implementation is available as a stand-alone open-source toolbox as well as a plugin for SnPM13 SnPM (2013), and is able to leverage multi-core architectures.

**Acknowledgements.** The authors thank Jia Xu for helping out with a preliminary implementation of the model. This work was supported in part by NIH R01 AG040396; NIH R01 EB022883; NIH R01 AG021155; NSF CAREER grant 1252725; NSF RI 1116584; UW ADRC P50 AG033514 and UW ICTR 1UL1RR025011. Hinrichs was supported by a CIBM post-doctoral fellowship at Wisconsin via NLM grant 2T15LM007359. Nichols and Maumet are supported by the Wellcome Trust (100309/Z/12/Z). The contents do not represent views of the Department of Veterans Affairs or the United States Government.

## 9. Supplementary Material

<http://pages.cs.wisc.edu/felipe/assets/pdf/rpt-nimg2017/SupplMaterial.pdf>

## References

A. Eklund, P. Dufort, D. F. S. L., 2013. Medical image processing on the gpu - past, present and future. *Medical Image Analysis* 17, 1073–1094.

- A. Eklund, M. Andersson, H. K., 2012. fmri analysis on the gpu-possibilities and challenges. *Computer Methods and Programs in Biomedicine* 105, 145–161.
- Arndt, S., Cizadlo, T., Andreasen, N., Heckel, D., Gold, S., O’Leary, D., 1996. Tests for comparing images based on randomization and permutation methods. *Journal of Cerebral Blood Flow and Metabolism* 16, 1271–1279.
- Balzano, L., Nowak, R., Recht, B., Sept 2010. Online identification and tracking of subspaces from highly incomplete information. In: *Communication, Control, and Computing (Allerton)*, 2010 48th Annual Allerton Conference on. pp. 704–711.
- Bennett, J., Lanning, S., 2007. The netflix prize. In: *Proceedings of KDD cup and workshop*. Vol. 2007. p. 35.
- Candès, E., Tao, T., 2010. The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inf. Theor.* 56 (5), 2053–2080.  
URL <http://dx.doi.org/10.1109/TIT.2010.2044061>
- Candès, E. J., Recht, B., 2009. Exact matrix completion via convex optimization. *Foundations of Computational mathematics* 9 (6), 717–772.
- Chandrasekaran, V., Sanghavi, S., Parrilo, P. A., Willsky, A. S., 2011. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization* 21 (2), 572–596.
- Cheverud, J. M., 2001. A simple correction for multiple comparisons in interval mapping genome scans. *Heredity* 87 (1), 52–58.
- Dahl, D. B., Newton, M. A., 2007. Multiple hypothesis testing by clustering treatment effects. *Journal of the American Statistical Association* 102 (478), 517–526.
- Edgington, E., 1969a. Approximate randomization tests. *The Journal of Psychology* 72 (2), 143–149.
- Edgington, E., 1969b. Statistical inference: The distribution-free approach.
- Edgington, E., Onghena, P., 2007. Randomization tests. CRC Press.
- Efron, B., Tibshirani, R., 2007. On testing the significance of sets of genes. *The annals of applied statistics*, 107–129.
- Eklund, A., Andersson, M., Knutsson, H., 2011. Fast random permutation tests enable objective evaluation of methods for single-subject fmri analysis. *International journal of biomedical imaging* 2011.
- Fazel, M., Hindi, H., Boyd, S., 2004. Rank minimization and applications in system theory. In: *American Control Conference, 2004. Proceedings of the 2004*. Vol. 4. IEEE, pp. 3273–3278.

- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., Frackowiak, R. S., 1994. Statistical parametric maps in functional imaging: a general linear approach. *Human brain mapping* 2 (4), 189–210.
- FSL, 2012. Fmrib software library (fsl). Available online at <http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/>.
- Gaonkar, B., Davatzikos, C., 2013. Analytic estimation of statistical significance maps for support vector machine based multi-variate image analysis and classification. *NeuroImage* 78, 270 – 283.  
URL <http://www.sciencedirect.com/science/article/pii/S1053811913003169>
- Gutierrez-Barragan, F., Ithapu, V., 2016. RapidPT. Available online at <https://github.com/felipegb94/RapidPT>.
- Halber, M., Herholz, K., Wienhard, K., Pawlik, G., Heiss, W.-D., 1997. Performance of a randomization test for single-subject 15o-water pet activation studies. *Journal of Cerebral Blood Flow & Metabolism* 17 (10), 1033–1039.
- He, J., Balzano, L., Szlam, A., June 2012. Incremental gradient on the grassmannian for online foreground and background separation in subsampled video. In: *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. pp. 1568–1575.
- Hinrichs, C., Ithapu, V., Sun, Q., Johnson, S., Singh, V., 2013. Speeding up permutation testing in neuroimaging. *Neural Information Processing Systems (NIPS)*.
- Holmes, A., Blair, R., Watson, J., Ford, I., 1996. Nonparametric analysis of statistic images from functional mapping experiments. *Journal on Cerebral Blood Flow and Metabolism* 16 (1), 7–22.
- Ji, H., Liu, C., Shen, Z., Xu, Y., 2010. Robust video denoising using low rank matrix completion. In: *CVPR. Citeseer*, pp. 1791–1798.
- Jockel, K.-H., 1984. *Computational Aspects of Monte Carlo Tests*. Physica-Verlag HD, Heidelberg.  
URL [http://dx.doi.org/10.1007/978-3-642-51883-6\\_25](http://dx.doi.org/10.1007/978-3-642-51883-6_25)
- Knijnenburg, T. A., Wessels, L. F., Reinders, M. J., Shmulevich, I., 2009. Fewer permutations, more accurate p-values. *Bioinformatics* 25 (12), i161–i168.
- Li, J., Ji, L., 2005. Adjusting multiple testing in multilocus analyses using the eigenvalues of a correlation matrix. *Heredity* 95 (3), 221–227.
- Nichols, T., Hayasaka, S., 2003. Controlling the familywise error rate in functional neuroimaging: a comparative review. *Statistical Methods in Medical Research* 12, 419–446.
- Nichols, T. E., Holmes, A. P., 2002. Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Map* 15 (1), 1–25.

- Recht, B., 2011. A simpler approach to matrix completion. *The Journal of Machine Learning Research* 12, 3413–3430.
- Recht, B., Fazel, M., Parrilo, P. A., 2010. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review* 52 (3), 471–501.
- SnPM, 2013. Statistical non parametric mapping toolbox (snpm). Available online at <http://www2.warwick.ac.uk/fac/sci/statistics/staff/academic-research/nichols/software/snpm>.
- SPM, 2012. Statistical parametric mapping toolbox (spm). Available online at <http://www.fil.ion.ucl.ac.uk/spm/>.
- Subramanian, A., Tamayo, P., Mootha, V. K., Mukherjee, S., Ebert, B. L., Gillette, M. A., Paulovich, A., Pomeroy, S. L., Golub, T. R., Lander, E. S., et al., 2005. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences* 102 (43), 15545–15550.
- Taylor, J., Worsley, K., 2008. Random fields of multivariate test statistics, with applications to shape analysis. *The Annals of Statistics* 36 (1), 1–27.
- Winkler, A., Ridgway, G., Douau, G., Nichols, T., Smith, S., 2016. Faster permutation inference in brain imaging. *NeuroImage* 141, 502–516.
- Winkler, A. M., Ridgway, G. R., Webster, M. A., Smith, S. M., Nichols, T. E., 2014. Permutation inference for the general linear model. *Neuroimage* 92, 381–397.
- Worsley, K., Evans, A., Marrett, S., Neelin, P., 1992. A three-dimensional statistical analysis for cbf activation studies in human brain. *Journal of Cerebral Blood Flow Metabolism* 12, 900–918.
- Worsley, K., Marrett, S., Neelin, P., Vandal, A., Friston, K., Evans, A., 1996. A unified statistical approach for determining significant signals in images of cerebral activation. *Human Brain Mapping* 4, 58–73.
- Y. Hochberg, A. T., 1987. *Multiple Comparison Procedures*, 1st Edition. Wiley.