A Conditional Randomization Test for Sparse Logistic Regression in High-Dimension

Binh T. Nguyen[1], **Bertrand Thirion**[2], **Sylvain Arlot**[3]

[1] *Télecom Paris, Palaiseau 91120, France*
tuanbinhs@gmail.com

[2] *Université Paris-Saclay, Inria, CEA, Palaiseau 91120, France*
bertrand.thirion@inria.fr

[3] *Université Paris-Saclay, CNRS, Inria, Laboratoire de mathématiques d'Orsay, 91405, Orsay, France*
sylvain.arlot@universite-paris-saclay.fr

Abstract: Identifying the relevant variables for a classification model with correct confidence levels is a central but difficult task in high-dimension. Despite the core role of sparse logistic regression in statistics and machine learning, it still lacks a good solution for accurate inference in the regime where the number of features p is as large as or larger than the number of samples n. Here we tackle this problem by improving the Conditional Randomization Test (CRT). The original CRT algorithm shows promise as a way to output p-values while making few assumptions on the distribution of the test statistics. As it comes with a prohibitive computational cost even in mildly high-dimensional problems, faster solutions based on distillation have been proposed. Yet, they rely on unrealistic hypotheses and result in low-power solutions. To improve this, we propose *CRT-logit*, an algorithm that combines a variable-distillation step and a decorrelation step that takes into account the geometry of ℓ_1-penalized logistic regression problem. We provide a theoretical analysis of this procedure, and demonstrate its effectiveness on simulations, along with experiments on large-scale brain-imaging and genomics datasets.

1. Introduction

Logistic regression is one of the most popular tools in modern applications of statistics and machine learning, partly due to its relative algorithmic simplicity. The method belongs to the class of *generalized linear models* that handle discrete outcomes, *i.e.* classification problems. Here, we focus on the binary classification problem, where one observation of the responses $y \in \{0, 1\}$ and the data vectors $\mathbf{x} \in \mathbb{R}^p$ follows the relationship:

$$\mathbb{P}(y = 1 \mid \mathbf{x}) = g(\mathbf{x}^T \boldsymbol{\beta}^0) = \frac{1}{1 + \exp(-\mathbf{x}^T \boldsymbol{\beta}^0)}, \tag{1}$$

where $g(x) = 1/(1 + \exp(-x))$ is the sigmoid function, and $\boldsymbol{\beta}^0$ the vector of true regression coefficients. In the classical setting, in which the number of samples n is greater than the number of features p, an estimate $\hat{\boldsymbol{\beta}}$ of the true signals $\boldsymbol{\beta}^0$ can be obtained using *maximum likelihood estimation* (MLE). The asymptotic behaviour and derivation of the test statistic, confidence intervals and p-values of the MLE have been well studied, *e.g.* in Cox and Hinkley (1979). The availability of p-values for the test statistics makes it possible to rely on multiple hypothesis testing, where one wants to test which variables have a non-zero effect on the outcome, *conditionally* to the remaining variables, *i.e.*

$$\text{(null)} \ \mathcal{H}_0^j : x_j \perp y \mid \mathbf{x}_{-j} \quad \text{vs.} \quad \text{(alternative)} \ \mathcal{H}_\alpha^j : x_j \not\perp y \mid \mathbf{x}_{-j},$$

for each feature $j \in [p] \stackrel{\text{def.}}{=} \{1, \ldots, p\}$ and $\mathbf{x}_{-j} \stackrel{\text{def.}}{=} \{x_1, x_2, \ldots, x_{j-1}, x_{j+1}, \ldots, x_p\}$. Equivalently, under the setting in Eq. (1), we have:

$$\text{(null)} \ \mathcal{H}_0^j : \beta_j^0 = 0 \quad \text{vs.} \quad \text{(alternative)} \ \mathcal{H}_\alpha^j : \beta_j^0 \neq 0.$$

Unfortunately, this line of analysis cannot be applied to the high-dimensional regime, where p is larger than n, as argued in Sur and Candès (2019); Yadlowsky et al. (2021); Zhao et al. (2022). These works show that in the regime where $\lim_{n,p \to \infty} n/p = \kappa$, the MLE estimator exists only when $\kappa > 2$. However, we note that this type of analysis is done *without* the addition of ℓ_1-regularization to the likelihood function, *i.e.* without using a *penalized estimator* to enforce sparsity.

Motivation Our focus in this paper is to do inference with statistical guarantees on high-dimensional sparse logistic regression, where p is larger or much larger than n. This setting is typical in modern applications of pattern recognition, *e.g.* in brain-imaging or genomics (Bach et al., 2012), with p as large as hundreds of thousands –compressible to thousands– but n stays at most few thousand. The family of methods
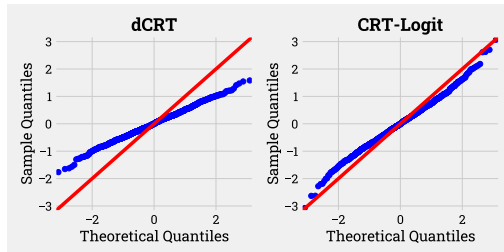

Figure 1: **QQ-Plot for 1000 samples of test-statistic of a null index for logistic regression**, with simulated data, $n = 200, p = 400$. *Left:* Statistics obtained from running Distilled-CRT, and *Right:* from our proposed algorithm. The empirical distribution of dCRT null-statistic strays far from theoretical distribution, which is standard normal, while empirical distribution of CRT-logit's null test score is much closer.

we consider is the *Conditional Randomization Test (CRT)* (Candès et al., 2018). CRT relies on generating multiple noisy copies of original variables to output empirical p-values in high-dimensional inference problems. However, prohibitive computational cost makes CRT impractical, as discussed at length in Candès et al. (2018); Tansey et al. (2022); Berrett et al. (2020); Liu et al. (2021). There have been several lines of research attempting to fix this problem, most notably the *distilled Conditional Randomization Test (dCRT)* (Liu et al., 2021). This work introduced a *distillation step* as a replacement for the randomized sampling step to compute the importance statistics (see Section 2 for more details). It provides a way to output p-values for multiple types of regression and classification problems, assuming convergence to Gaussian of the test statistic in large-sample regime. Yet, as shown in the left panel of Figure 1, the originally proposed dCRT test-statistic for logistic regression does not behave as well as intended. *In particular, its null distribution deviates markedly from standard normal in high-dimension whenever $n/p \leq 1$.*

Contribution We propose a correction for the dCRT, inspired by the decorrelation method presented in Ning and Liu (2017). The decorrelation step makes the null-distribution of the test statistics much closer to standard normal, as shown on the right panel of Figure 1, and thus increases the statistical power of the method. We provide asymptotic analysis of this method, which shows that CRT-logit produces standard normal test-statistic in large-sample regime. In addition, we validate the high performance of CRT-logit on large-scale brain-imaging and genetics datasets, thus showing its usefulness in practical applications.

Related works The closest cousin of the Conditional Randomization Test is Knock-off Filter (Barber and Candès, 2015; Candès et al., 2018), a recent breakthrough in the False Discovery Rate (FDR) control literature. It relies on the creation of additional noisy features, called knockoffs, to calculate variable-importance statistics. Another extension of vanilla CRT is the Holdout Randomization Test (HRT) (Tansey et al., 2022). While still requiring multiple sampling of noisy variables, HRT solves the computational issue of original CRT by doing heavy model fitting only once on one part of the dataset, and test statistics calculation on the other part, without refitting the model. However, this method relies on sample-splitting, hence inherently suffers from a loss of statistical power. A parallel line of work has introduced the Conditional Permutation Test (CPT) (Berrett et al., 2020), a non-parametric alternative to CRT that relies on a random shuffling mechanism applied to original variables, instead of multiple sampling of new variables. This potentially makes CPT more robust to model mis-specification. Yad-lowsky et al. (2021) recently proposed a method called SLOE, which adapts the analysis

of Zhao et al. (2022), but in the regime different from what we are considering, where $\lim_{n,p\to\infty} n/p \to \kappa \in (1,2)$, and more importantly without sparsity-inducing penalty. On a separate note, we notice the similarity of dCRT (Liu et al., 2021) with debiased Lasso (Javanmard and Montanari, 2014; van de Geer et al., 2014; Zhang and Zhang, 2014). This line of work proposed a debiasing formula for the estimator, which makes the asymptotic distribution of $(\hat{\boldsymbol{\beta}}^{\mathrm{LASSO}} - \boldsymbol{\beta}^0)$ standard normal, so that one can compute the test statistic and p-value associated with each variable.

2. Background

Notation We denote matrices, vectors, scalars and sets by bold uppercase, bold lowercase, script lowercase , and calligraphic letters, respectively, *e.g.* $\mathbf{X}$, $\mathbf{x}$, x, and $\mathcal{X}$. The i-th row of a matrix $\mathbf{X}$ will be denoted $\mathbf{X}_{i,*}$, the j-th column $\mathbf{X}_{*,j}$ and the (i,j)-th element $\mathbf{X}_{i,j}$. For any natural number p, we denote the set $[p] \overset{\text{def.}}{=} \{1,\ldots,p\}$. For each $\mathbf{x} \in \mathbb{R}^p$ and $j \in [p]$, we denote $\mathbf{x}_{-j} \overset{\text{def.}}{=} \{x_1, x_2, \ldots, x_{j-1}, x_{j+1}, \ldots, x_p\}$ a $p-1$ dimensional observation after removing the j-th variable. Correspondingly, $\mathbf{X}_{-j}$ is the data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ with column $\mathbf{X}_{*,j}$ removed. The cumulative distribution function (CDF) of the standard Gaussian distribution will be denoted $\Phi(\cdot)$. The indicator function of a random event $\mathcal{A}$ will be denoted $\mathbf{1}_{\mathcal{A}}$. For any two positive sequences x_n and y_n, we write $x_n \asymp y_n$ if $cy_n \leq x \leq Cy_n$ for all n, for some positive constants c and C. For a vector $\mathbf{x}$, $\|\mathbf{x}\|_p$ denotes its ℓ_p norm. For a function $f : \mathbb{R}^p \to \mathbb{R}$, $\nabla_j f$ denotes its gradient *w.r.t.* the j-th variable, for $j \in [p]$.

Problem setting We consider exclusively binary classification, where the response vector is denoted $\mathbf{y} \in \{0,1\}^n$ and the data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ consists of n p-dimensional samples. Throughout the paper, we assume the data $\{\mathbf{X}_{i,*}\}_{i=1}^n$ are *i.i.d.* and follow a distribution with zero mean and population covariance matrix $\boldsymbol{\Sigma}$. Moreover, we assume that $\mathbf{X}_{i,*}$ and $\mathbf{y}_i$ follow the logistic relationship in Eq. (1). We denote the support set $\mathcal{S} \overset{\text{def.}}{=} \{j \in [p] : \boldsymbol{\beta}_j^0 \neq 0\}$ and assume that it is sparse, *i.e.* $\mathrm{card}(\mathcal{S}) = s^* \ll p$, where card denotes the cardinality of a set. Furthermore, $\hat{\mathcal{S}} \overset{\text{def.}}{=} \{j \in [p] : \hat{\beta}_j \neq 0\}$ indicates an estimation of $\mathcal{S}$, where $\hat{\beta}_j$ is an estimate of the true signal $\boldsymbol{\beta}_j^0$. We try to obtain it through a ℓ_1-penalized logistic estimator:

$$\hat{\boldsymbol{\beta}}_\lambda = \underset{\boldsymbol{\beta} \in \mathbb{R}^p}{\mathrm{argmin}}\ \ell(\boldsymbol{\beta}) + \lambda\|\boldsymbol{\beta}\|_1\ ,\ \text{with } \ell(\boldsymbol{\beta}) = -\frac{1}{n}\sum_{i=1}^n \{y_i(\mathbf{X}_{i,*}\boldsymbol{\beta}) - \log\left[1 + \exp(\mathbf{X}_{i,*}\boldsymbol{\beta})\right]\}\ .$$
(2)

We denote $\mathbf{I} \overset{\text{def.}}{=} \mathbb{E}_{\boldsymbol{\beta}^0}[\nabla^2 \ell(\boldsymbol{\beta}^0)]$ the Fisher information matrix, and $\mathbf{I}_{j|-j}$ the partial Fisher information, defined by $\mathbf{I}_{j|-j} \overset{\text{def.}}{=} \mathbb{E}[\nabla_{jj}^2 \ell(\boldsymbol{\beta}^0) - [\nabla_{j,-j}^2 \ell(\boldsymbol{\beta}^0)]^\top [\nabla_{-j,-j}^2 \ell(\boldsymbol{\beta}^0)]^{-1} \nabla_{-j,j}^2 \ell(\boldsymbol{\beta}^0)] = \mathbf{I}_{jj} - \mathbf{I}_{j,-j}\mathbf{I}_{-j,-j}^{-1}\mathbf{I}_{-j,j}$, where $\mathbf{I}_{j,-j}$ is the row-vector made with the jth-row and the columns corresponding to $\boldsymbol{\beta}_{-j}$, $\mathbf{I}_{-j,-j}$ the sub-matrix of $\mathbf{I}$ made with the rows and columns corresponding to $\boldsymbol{\beta}_{-j}$. This quantity, defined following Cox and Hinkley (1979, pp. 323), plays an important role in our proposed method, detailed in Section 3.

Statistical control with False Discovery Rate To quantify statistical errors, we consider the *False Discovery Rate*, introduced in Benjamini and Hochberg (1995). Given an estimate of the support $\hat{\mathcal{S}}$, the false discovery proportion (FDP) is the ratio of the number of selected features that do not belong to the true support $\mathcal{S}$, divided by the total number of selected features. The False Discovery Rate is the expectation of the FDP:

$$\mathrm{FDP}(\hat{\mathcal{S}}) = \frac{\mathrm{card}(\{j : j \in \hat{\mathcal{S}}, j \notin S\})}{\mathrm{card}(\hat{\mathcal{S}}) \vee 1} \qquad \text{and} \qquad \mathrm{FDR}(\hat{\mathcal{S}}) = \mathbb{E}[\mathrm{FDP}(\hat{\mathcal{S}})].$$

Conditional Randomization Test (CRT) and Distillation CRT (dCRT) The concept of Conditional Randomization Test was originally proposed in the model-X

knockoff paper (Candès et al., 2018) as a way to output valid empirical p-values using knockoff variables. The principle of the knockoff filter is first to sample noisy copies $\tilde{\mathbf{X}}_{*,j}$ of variable $\mathbf{X}_{*,j}$, given a known sampling mechanism $P_{j\ |-j}$. One advantage of the knockoff filter is that no specific assumption is placed on the distribution of the inferred test statistic. However, this means that, in general, there is no mechanism to derive p-values from the knockoff statistic. This motivates the introduction of CRT, which requires running high-dimensional inference for each variable j B times. However, the computation cost of CRT is prohibitive when p grows large: assuming that we use the Lasso program with coordinate descent to compute T_j^{CRT}, its runtime would be $\mathcal{O}(Bp^4)$ (Hastie et al., 2009, pp. 93). Moreover, CRT requires decently large B to make the empirical distribution of the p-values smooth enough. Reducing the computational cost of CRT is the main motivation of several works (Berrett et al., 2020; Liu et al., 2021; Tansey et al., 2022). One of them is the introduction of distillation-CRT (dCRT) by Liu et al. (2021). The main appeal of this method is that it can output p-values analytically, therefore bypassing the multiple knockoffs sampling steps used to infer on each variable, and leads to a reasonable reduction of the computation cost.

Distillation operation The key addition of dCRT is the distillation operation: for each variable j, we want to distill all the conditional information of the remaining variables $\mathbf{X}_{-j}$ to $\mathbf{x}_j$ and to $\mathbf{y}$ via least-squares minimization with ℓ_1-regularization to enforce sparsity. For each variable j, we first solve the lasso problem by regressing $\mathbf{X}_{*,j}$ on $\mathbf{X}_{-j}$,

$$\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} = \underset{\boldsymbol{\beta}\in\mathbb{R}^{p-1}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{X}_{*,j} - \mathbf{X}_{-j}\boldsymbol{\beta}\|_2^2 + \lambda_{dx}\|\boldsymbol{\beta}\|_1. \tag{3}$$

For distillation of variable j and the binary response $\mathbf{y}$ with logistic relationship, Liu et al. (2021) briefly suggested to solve a penalized estimation problem, similar to Eq. (2):

$$\hat{\boldsymbol{\beta}}^{d_y,j} = \underset{\boldsymbol{\beta}\in\mathbb{R}^{p-1}}{\operatorname{argmin}} -\frac{1}{n}\sum_{i=1}^n \left\{ y_i(\mathbf{X}_{i,-j}^\top\boldsymbol{\beta}) - \log\left[1 + \exp(\mathbf{X}_{i,-j}^T\boldsymbol{\beta})\right]\right\} + \lambda\|\boldsymbol{\beta}\|_1. \tag{4}$$

Then, a test statistic is calculated for each $j = 1,\ldots,p$:

$$T_j = \sqrt{n}\, \frac{\langle\mathbf{y} - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d_y,j}, \mathbf{X}_{*,j} - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}}\rangle}{\|\mathbf{y} - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d_y,j}\|_2\|\mathbf{x}_j - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}}\|_2}, \tag{5}$$

which, under the null hypothesis, and more importantly, assuming linear relationship between $\mathbf{X}_{i,*}$ and $\mathbf{y}$, follows standard normal distribution asymptotically, conditional to $\mathbf{y}$ and $\mathbf{X}_{-j}$. It then follows that we can output a p-value for each variable j by taking $\hat{p}_j = 2\left[1 - \Phi\left(T_j\right)\right]$.

However, the formulation of test statistics in Eq. (5) is not truly satisfactory in the setting of sparse logistic regression. More specifically, both the calculation of regression residuals $\mathbf{y} - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d_y,j}$ and test statistics T_j *do not take into account the non-linear relationship* between $\mathbf{X}$ and the binary response $\mathbf{y}$. The first row of Figure 6 plots the qq-plot of the test statistics T_j for logistic regression, which shows that even in the classical regime where $n > p$, its distribution is far from standard normal.

3. Decorrelating Test-Statistics for High-Dimensional Logistic Regression

As we have elaborated, the formulation of dCRT is not well-suited for problems other than penalized least-squares regression. We therefore propose an adaptation of dCRT in the case of logistic regression, inspired by the classical work of Cox and Hinkley (1979) and by Ning and Liu (2017). First, note that when testing $H_0^j : \beta_j^0 = 0$ under the case where $n > p$, we have the classical Rao's test statistic, defined by

$$T_j^{\mathrm{Rao}} = \sqrt{n}\,\hat{\mathbf{I}}_{j|-j}^{-1/2}\,\nabla_j\ell(0,\hat{\boldsymbol{\beta}}_{-j}), \tag{6}$$

where $\nabla_j\ell(0,\hat{\boldsymbol{\beta}}_{-j}) \overset{\text{def.}}{=} \nabla_{\beta_j}\ell(\beta_j,\hat{\boldsymbol{\beta}}_{-j})\Big|_{\beta_j=0}$ is the Fisher score. Here $\hat{\boldsymbol{\beta}}_{-j} \overset{\text{def.}}{=} \operatorname{argmin}_{\boldsymbol{\beta}_{-j}\in\mathbb{R}^{p-1}}\ell(\beta_j,\boldsymbol{\beta}_{-j})$ is the constrained maximum-likelihood estimator of $\boldsymbol{\beta}_{-j}$ with fixed β_j, and $\hat{\mathbf{I}}_{j|-j}$ is a

consistent estimator of the partial Fisher information $\mathbf{I}_{j|-j}$. The appearance of the term $\hat{\mathbf{I}}_{j|-j}^{-1/2}$ is due to the fact that under the null hypothesis H_0^j, we have, by Cox and Hinkley (1979, Chapter 9), and by Rao (1948),

$$\sqrt{n}\nabla_j\ell(0,\hat{\boldsymbol{\beta}}_{-j}) \xrightarrow[n\to\infty]{(d)} \mathcal{N}(0,\mathbf{I}_{j|-j}),$$

which makes the asymptotic distribution of T_j^{Rao} standard normal. However, in the high-dimension case, where $n < p$, we do not reach this convergence in distribution. To see this, consider the Taylor expansion of the Fisher score of variable j around any given estimator $\widetilde{\boldsymbol{\beta}}_{-j}$ of the true $\boldsymbol{\beta}_{-j}^0$:

$$\nabla_j\ell(0,\widetilde{\boldsymbol{\beta}}_{-j}) = \nabla_j\ell(0,\boldsymbol{\beta}_{-j}^0) + \nabla_{j,-j}^2\ell(0,\boldsymbol{\beta}_{-j}^0)(\widetilde{\boldsymbol{\beta}}_{-j} - \boldsymbol{\beta}_{-j}^0) + \mathcal{O}\left((\widetilde{\boldsymbol{\beta}}_{-j} - \boldsymbol{\beta}_{-j}^0)^2\right) \quad (7)$$

On the right-hand side, the first term converges weakly to a normal distribution due to the Central Limit Theorem, the remainder term becomes negligible using ℓ_1 regularization to induce sparsity, but the second term does not, due to estimation bias and sparsity effect of $\widetilde{\boldsymbol{\beta}}_{-j}$ (Fu and Knight, 2000).

Adapting distillation operation for sparse logistic regression Fortunately, Eq. (7) suggests that for each variable j, we can *debias* the Fisher score by correcting the impact of other terms. In particular, for each variable j, we replace the Fisher score by

$$\nabla_j\ell(0,\boldsymbol{\beta}_{-j}) - \mathbf{I}_{j,-j}\mathbf{I}_{-j,-j}^{-1}\nabla_{-j}\ell(0,\boldsymbol{\beta}_{-j}) . \quad (8)$$

The inversion of the large matrix $\mathbf{I}_{-j,-j} \in \mathbb{R}^{(p-1)\times(p-1)}$ is computationally prohibitive, but we can estimate the term $\mathbf{I}_{j,-j}\mathbf{I}_{-j,-j}^{-1}$ straightforwardly by solving

$$\hat{\mathbf{w}}^j = \underset{\mathbf{w}\in\mathbb{R}^{p-1}}{\operatorname{argmin}} \frac{1}{2n}\sum_{i=1}^n \left[\nabla_{j,-j}^2\ell_i(\hat{\boldsymbol{\beta}}) - \mathbf{w}^T\nabla_{-j,-j}^2\ell_i(\hat{\boldsymbol{\beta}})\right]^2 + \lambda\|\mathbf{w}\|_1, \quad (9)$$

for each variable j, where $\hat{\boldsymbol{\beta}}$ is given with Eq. (2). Moreover, since we have the closed-form of the derivatives of the logistic loss $\ell(\hat{\boldsymbol{\beta}})$, a simple derivation from Eq. (9) suggests the following x_j-distillation operation, *adapted for logistic regression*:

$$\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} = \underset{\boldsymbol{\beta}\in\mathbb{R}^{p-1}}{\operatorname{argmin}} \frac{1}{n}\sum_{i=1}^n g''(\mathbf{X}_{i,*}\hat{\boldsymbol{\beta}})(\mathbf{X}_{i,j} - \mathbf{X}_{i,-j}\boldsymbol{\beta})^2 + \lambda_{dx}\|\boldsymbol{\beta}\|_1, \quad (10)$$

where the extra term (second-order derivative of the sigmoid function) $g''(\mathbf{X}_{i,*}\hat{\boldsymbol{\beta}}) = \frac{\exp(\mathbf{X}_{i,*}\hat{\boldsymbol{\beta}})}{[1+\exp(\mathbf{X}_{i,*}\hat{\boldsymbol{\beta}})]^2}$ appears from differentiating twice the loss function $\ell(\hat{\boldsymbol{\beta}})$, and $\hat{\boldsymbol{\beta}} = \hat{\boldsymbol{\beta}}_\lambda$ is defined in Eq. (2). On the other hand, we can obtain $\hat{\boldsymbol{\beta}}_j^{d_y,j}$ from $\hat{\boldsymbol{\beta}}$ by simply omitting the j-th coefficient from it, *i.e.*

$$\hat{\boldsymbol{\beta}}^{d_y,j} \overset{\mathrm{def.}}{=} (\hat{\beta}_1, \hat{\beta}_2, \ldots, \hat{\beta}_{j-1}, \hat{\beta}_{j+1}, \ldots, \hat{\beta}_p) .$$

Finally, the equation for decorrelated test score, adapted from both Eq. (5) and (6), reads

$$T_j^{\mathrm{decorr}} = -\frac{1}{\sqrt{n}}\,\hat{\mathbf{I}}_{j|-j}^{-1/2}\,\sum_{i=1}^n \left[y_i - g(\mathbf{X}_{i,-j}\hat{\boldsymbol{\beta}}^{d_y,j})\right]\left[\mathbf{X}_{i,j} - \mathbf{X}_{i,-j}\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}}\right], \quad (11)$$

where the formula for empirical partial Fisher information is $\hat{\mathbf{I}}_{j|-j} = n^{-1}\sum_{i=1}^n g''(\mathbf{X}_{i,*}\hat{\boldsymbol{\beta}})(\mathbf{X}_{i,j} - \mathbf{X}_{i,-j}\,\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}})\,\mathbf{X}_{i,j}$. A summary of the full procedure, which we call CRT-logit, can be found in Algorithm 1. Notice that the runtime of CRT-logit is the same as dCRT, which means in general slower than KO and HRT. To speedup inference time, we introduce a variable-screening step that eliminates potentially unimportant variables before distillation, similar to dCRT. We provide empirical benchmark of runtime of each method in Section 4.5.

Setting ℓ_1-regularization parameter λ and λ_{dx} In general, we advise to use cross-validation for obtaining $\hat{\boldsymbol{\beta}}_\lambda$ in Eq. (2) and for $\mathbf{X}_{*,j}$-distillation operator, as defined by Eq. (10). This is inline with the theoretical argument for dCRT (Liu et al., 2021, Lemma 1 and Theorem 3). However, we also observe empirically that choosing the ℓ_1-regularization parameters of the distillation step can strongly affect how variables are selected by CRT-logit. We provide more details Appendix C, and leave further theoretical investigations of this phenomenon for future work.

Algorithm 1: CRT-logit

---

1 **INPUT** design matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$, reponses $\mathbf{y} \in \mathbb{R}^n$

2 **OUTPUT** vector of p-values $\{p_j\}_{j=1}^p$;

3 $\hat{\boldsymbol{\beta}}_\lambda \leftarrow$ `solve_sparse_logistic_cv`$(\mathbf{X}, \mathbf{y})$ // Using Eq. (2)

4 $\hat{\mathcal{S}}^{\text{SCREENING}} \leftarrow \{j : j \in [p], \hat{\beta}_j \neq 0\}$

5 **for** $j \in \hat{\mathcal{S}}^{SCREENING}$ **do**

6 $\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} \leftarrow$ `solve_scaled_lasso_cv`$(\mathbf{X}_{*,j}, \mathbf{X}_{*,-j})$ // Using Eq. (10)

7 $\hat{\boldsymbol{\beta}}^{d_y,j} \leftarrow (\hat{\beta}_1, \hat{\beta}_2, \ldots, \hat{\beta}_{j-1}, \hat{\beta}_{j+1}, \ldots, \hat{\beta}_p)$

8 $T_j^{\text{decorr}} \leftarrow$ `decorrelated_test_score`$(j, \mathbf{X}, \mathbf{y}, \hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}}, \hat{\boldsymbol{\beta}}^{d_y,j})$ // Using Eq (11)

9 $\hat{p}_j \leftarrow 2[1 - |\Phi\left(T_j^{\text{decorr}}\right)|]$

10 **end**

11 **for** $j \notin \hat{\mathcal{S}}^{SCREENING}$ **do**

12 $\hat{p}_j = 1$

13 **end**

Asymptotic analysis of the Decorrelated Test Statistic We now provide theoretical analysis of CRT-logit in large-sample regime. All the proofs can be found in Appendix A. Without writing it explicitly, in our analysis, we consider $p = p(n)$, and the following assumption.

Assumption 3.1 (Regularity conditions)**.** *Assume that*

(A1) $\lambda_{\min}(\mathbf{I}) \geq \kappa^2$ *for some constant* $\kappa > 0$.

(A2) *Sparsity of* $\boldsymbol{\beta}^0$ *and* $\mathbf{w}^{0,j}$, *with* $\mathbf{w}^{0,j}$ *the ground truth weights for the distillation of* $\mathbf{x}_j$ *in Eq.* (10)*:* $|\mathcal{S}| = s^*$ *and* $\|\mathbf{w}^{0,j}\|_0 = s'$ *with* $s^* = o\left(n^{1/2}/\log(p)\right)$ *and* $s' = o\left(n^{1/2}/\log(p)\right)$.

(A3) *For all* $i \in [n]$, $\mathbf{X}_{i,*}$ *and* $(-y_i + g'(\mathbf{X}_{i,*}\beta))$ *are sub-exponential random variables, and* $|\mathbf{X}_{i,-j}\mathbf{w}^{0,j}| \leq K$ *almost surely, for some constant* K.

We then have the following result, that states that the asymptotic distribution of the decorrelated test scores is standard normal.

Theorem 3.1. *Let* $j \in [p]$, *and let* T_j^{decorr} *be defined as in Eq.* (11)*, with* $\lambda \asymp \lambda_{dx} \asymp \sqrt{n^{-1}\log(p)}$. *Then, if Assumption 3.1 holds true, under the null hypothesis* $\mathcal{H}_0^j : \beta_j^0 = 0$, *we have*

$$\forall t \in \mathbb{R}, \qquad \lim_{n \to \infty} |\mathbb{P}_{\beta^0}(T_j^{decorr} \leq t) - \Phi(t)| = 0,$$

where $\Phi(\cdot)$ *is the CDF of the standard Gaussian distribution. Moreover, for each* $j \in [p]$, *if we define* $\hat{p}_j \overset{def.}{=} 2\left[1 - \Phi(T_j)\right]$, *i.e.* $\hat{p}_j$ *is the output of Algorithm 1, then, under the null hypothesis* $\mathcal{H}_0^j : \beta_j^0 = 0$, *we have*

$$\limsup_{n \to \infty} \mathbb{P}_{\beta^0}(\hat{p}_j \leq t) \leq t \quad \text{for all } t \in [0, 1],$$

that is, the p-values output by Algorithm 1 are valid asymptotically.

FDR control with CRT-logit We now state the second main result, which establishes that the FDR of the test is controlled when using Benjamini-Yekutieli procedure (Benjamini and Yekutieli, 2001) with the p-values output from Algorithm 1, assuming that the number of tests p is fixed.

Theorem 3.2. *Under Assumptions 3.1 and logistic model defined in Eq (1), with $\lambda \asymp \lambda_{dx} \asymp \sqrt{n^{-1}\log(p)}$, assume $n^{-1/2}(s' \vee s^*)\log(p) = o(1)$, and assume the number of tests p is fixed. Let $\alpha \in (0,1)$ and $\widehat{\mathcal{S}}_{BY\text{-}CRT}$ be given by applying following the Benjamini-Yekutieli FDR-controlling procedure to the CRT-logit p-values $\{\widehat{p}_j\}_{j \in [p]}$, output from Algo.1. Then, we have*

$$\limsup_{n \to \infty} \mathbb{E}\left[\frac{\operatorname{card}(\widehat{\mathcal{S}}_{BY\text{-}CRT} \cap \mathcal{S}^c)}{\operatorname{card}(\widehat{\mathcal{S}}_{BY\text{-}CRT}) \vee 1}\right] \leq \alpha \,.$$

Remark 3.1. *Assumption 3.1 is also assumed in Ning and Liu (2017); van de Geer et al. (2014), which also provide a detailed discussion of this regularity assumption in generalized linear models. This assumption, in turn, is built on the regularity assumption in the classic work of Cox and Hinkley (1979, Chapter 9) to establish asymptotic normality of Rao's test statistic. Theorem 3.1 is an adaptation of Ning and Liu (2017, Theorem 3.1), specialized for the case of sparse logistic regression and the p-values output from CRT-logit.*

4. Empirical Results

We provide benchmarks of the proposed CRT-logit algorithm along with most other methods mentioned in the introduction, in particular: model-X Knockoff (KO) (Candès et al., 2018), Debiased Lasso (dLasso) (Zhang and Zhang, 2014; Javanmard and Montanari, 2014), original CRT with 1000 samplings (Candès et al., 2018), Holdout Randomization Test with 5000 samplings (Tansey et al., 2022), and Lasso-Distillation CRT (dCRT) (Liu et al., 2021). We did not include SLOE (Yadlowsky et al., 2021) and CPT (Berrett et al., 2020), as the provided open-source implementation are particularly unstable and do not fit in the sparse-regression setting (for SLOE), or implementation are not available (for CPT).

Remark 4.1. *As a slight caveat, in the simulated and semi-realistic experiment sections (Sections 4.1, 4.2 and 4.3), we introduce an additional noise term to the logistic relationship of Eq. (1):*

$$\mathbb{P}(y_i = 1 \mid \mathbf{x}_i) = g(\mathbf{x}_i^T \boldsymbol{\beta}^0 + \sigma \xi_i) \,, \tag{12}$$

where $\xi_i \sim \mathcal{N}(0,1)$ is a Gaussian noise and $\sigma > 0$ the noise magnitude. The formula in Eq. (12) has been used in previous works, e.g. Bzdok et al. (2015). There is a clear justification to this: in most of the applications we consider, data are collected with measurement errors. In the case of brain-imaging, for example, recording the brain signal of the human subjects by scanners often includes noise caused either from the machine, or from the movement of the subjects, as elaborated by Lindquist (2008). Moreover, in general, this setting corresponds to a model mis-specification, which the CRT-logit is robust to under Assumption 3.1, following the arguments as in Ning and Liu (2017, Section 5).

Remark 4.2. *We use Benjamini-Hochberg step-up procedure (Benjamini and Hochberg, 1995) to control FDR with the p-values in all the empirical experiments in Section 4.2 and App. 4.4, as we observe empirically the FDR is usually controlled with this procedure, without compromising power with the conservative BY bound.*

4.1. Effectiveness of the decorrelation step

To show how decorrelating the test statistics helps, we set up a simulation with matrix $\mathbf{X}$ of $p = 400$ features and vary the number of samples $n \in \{200, 400, 800, 4000\}$. The binary response vector $\mathbf{y}$ is created following Eq. (12), and the design matrix $\mathbf{X}$ is sampled from a multivariate normal distribution with zero mean, while the covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{p \times p}$ is a symmetric Toeplitz matrix, where the parameter $\rho \in (0,1)$ controls the correlation between neighboring features: correlation decreases quickly when the distance between feature indices increases. The true signal $\boldsymbol{\beta}^0$ is picked with a sparsity

parameter $\kappa = s^*/p$ that controls the proportion of non-zero elements with magnitude 2.0, *i.e.* $\beta_j = 2.0$ for all $j \in \mathcal{S}$. For the specific purpose of this experiment, non-zero indices of $\mathcal{S}$ are kept fixed. The noise $\boldsymbol{\xi}$ is *i.i.d.* normal $\mathcal{N}(0, \mathbf{Id}_n)$ with magnitude $\sigma = \|\mathbf{X}\boldsymbol{\beta}^0\|_2/(\sqrt{n}\ \text{SNR})$, controlled by the SNR parameter. In short, the three main parameters controlling this simulation are correlation ρ, sparsity degree κ and signal-to-noise ratio SNR. We generate randomly 1000 datasets, and run dCRT and CRT-logit algorithm to obtain one sample of test statistics $\{T_j\}_{j=1}^p$ and $\{T_j^{\text{decorr}}\}_{j=1}^p$. Then, we pick 1000 samples of one null test statistic T_j and T_j^{decorr}, defined in Eq. (5) and (11), respectively, and plot the qq-plot of their empirical quantile versus the standard normal quantile. From the results in Figure 6, we observe that the empirical null distribution of
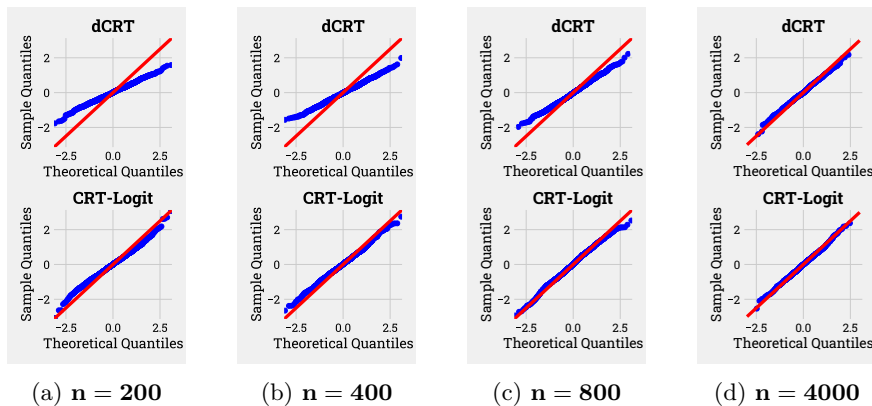

(a) **n = 200** (b) **n = 400** (c) **n = 800** (d) **n = 4000**

Figure 2: **QQ-Plot for one null CRT statistic for logistic regression, with varying number of samples and a fixed number of variables** $p = 400$. The theoretical quantiles are obtained from a standard Gaussian distribution. The decorrelation step makes the empirical null distribution of the null statistics much closer to standard Gaussian. Parameters: SNR $= 3.0, \rho = 0.4$, sparsity $= 0.06$. *Upper row: Distilled-CRT statistic defined by Eq. (5). Bottom row: CRT-logit, with decorreleated test score defined by Eq. (11) **(ours)**.*

the test statistic is *much closer to a standard normal when adding the decorrelation step.* In particular, when the sample size n increases to 400, the decorrelated test statistic has empirical quantiles almost inline with the theoretical quantiles of the standard normal distribution, while dCRT test score strays away from the 45-degree line. Again, we emphasize that the normality of T_j is crucial for the p-values calculation. This outlines the importance of the decorrelating step on T_j.

4.2. High-dimensional scenario with varying simulation parameters

To see how each algorithm performs under different settings, we follow the same simulation scenario as in Sec. 4.1, but vary each of the three simulation parameters, while keeping the others unchanged at default value of SNR $= 2.0, \rho = 0.5, \kappa = 0.04$. We target a control of FDR at level 0.1, using Benjamini-Hochberg procedure. Results in Figure 3 show that CRT-logit is the most powerful method while still controlling the FDR. Moreover, in the presence of higher correlations between nearby variables ($\rho > 0.6$), other methods suffer a drop in average power, but this is not as severe for CRT-logit. The original CRT, in general, is conservative. We believe that this is due to using only $B = 500$ samplings to generate empirical p-values for the two methods, due to prohibitive average runtime of the algorithm with larger B (which we provide in Section 4.5). For HRT, the conservativeness is expected, due to the usage of only half of the sample for test-statistics calculation – even though the number of samplings is bigger than original CRT ($B = 5000$). We note that, perhaps surprisingly, the debiased lasso (`cdlasso`) is the most conservative. It controls FDR well in all settings. This might be due to the fact that `dlasso` also relies on the choice of the ℓ_1-regularization λ in the nodewise Lasso operation, similar to the $\mathbf{X}_{*,j}$-distillation of dCRT, as noted in Section 1. What makes

the difference is that instead of using cross-validation for setting λ for each variable j, a *fixed* value of $\lambda = 10^{-2}\lambda_{max}$ is used in the implementation of `dlasso`. We strongly suspect this fixed value is not optimal, which makes the procedure powerless.
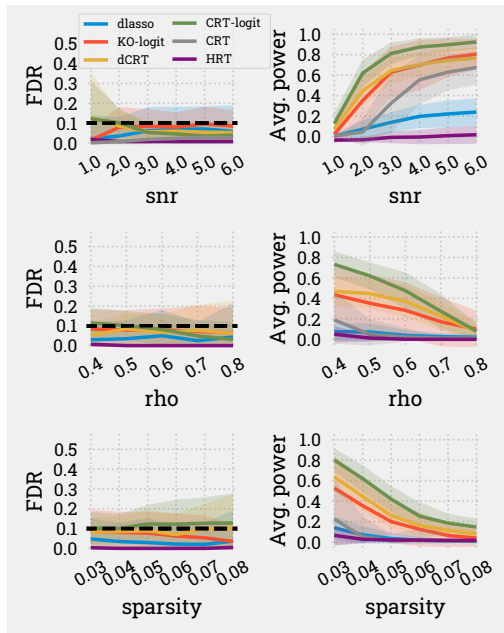

Figure 3: **FDR/Average Power of 100 runs of simulations across varying parameters in high-dimensional settings**. Default parameter: $n = 400, p = 600, \mathrm{SNR} = 2.0, \rho = 0.5, \kappa = 0.04$. FDR is controlled at level $\alpha = 0.1$. Methods: Debiased Lasso (`dlasso`), model-X Knockoff (`KO-logit`), original CRT (`CRT`), HRT (`HRT`), dCRT (`dCRT`), and our version of CRT (dark green line – `CRT-logit`).

4.3. Application: large-scale analysis on brain-imaging dataset

Description   The Human Connectome Project dataset (HCP) is a collection of brain imaging data on healthy young adult subjects with age ranging from 22 to 35. More specifically, the input $\mathbf{X}$ is a set of 2mm statistical maps of 400 subjects across 8 cognitive tasks. These are called z-maps, as the data follow a standard normal distribution under the null hypothesis. Each task in turn features 2 different contrasts, which effectively form binary responses $\mathbf{y} \in \{0, 1\}^n$. We propose to fit $\mathbf{y}$ through distributed brain signals and identify relevant brain locations. The setting is high-dimensional with $n = 800$ samples, corresponding to 400 subjects, while the total number of variables is $p \approx 200,000$ brain voxels. Following Nguyen et al. (2019); Chevalier et al. (2021), we use a hierarchical clustering scheme to group the variables into $C = 1000$ spatially connected clusters. We provide details of the pre-processing step in Appendix F.

Creating semi-realistic ground-truth and response labels   Since there is no ground truth for this dataset, we create synthetic true signals by fitting the data $\mathbf{X}$ and response $\mathbf{y}$ with an ℓ_1-penalized logistic classifier. In other words, the estimator $\hat{\boldsymbol{\beta}}^{\texttt{logreg}}$ will serve as true regression coefficients for each task. Then, to avoid bias in simulating label $\hat{\mathbf{y}}$, the z-maps matrix $\mathbf{X}$ of one task are used in conjunction with the discriminative pattern map $\hat{\boldsymbol{\beta}}^{\texttt{logreg}}$ from the next task in the following order: `relational`, `gambling`, `emotion`, `social`. For instance, we use $\hat{\boldsymbol{\beta}}^{\texttt{logreg}}$ of `gambling` with z-maps data matrix of `relational`, *i.e.* for all $i = 1, \ldots, n$, given $\mathbf{x}_{i,\texttt{relational}}$,

$$\hat{y}_i \sim \mathrm{Bern}\left\{ g(\mathbf{x}_{i,\texttt{relational}}^\top \, \hat{\boldsymbol{\beta}}_{\texttt{gambling}}^{\texttt{logreg}} + \sigma \xi_i) \right\}, \qquad (13)$$

where $\mathrm{Bern}(a)$ is a Bernoulli probability mass function that takes a value 1 with probability a, σ is a noise magnitude and ξ_i is a standard normal noise. Finally, we apply

all inference algorithms on the semi-synthetic data $(\mathbf{X}, \hat{\mathbf{y}})$, and we evaluate their performance using the ground-truth $\hat{\boldsymbol{\beta}}^{\mathrm{logreg}}$. This simulation setting is similar to Chevalier et al. (2021), except that here we consider a classification and not a regression problem. It allows us to calculate the False Discovery Rate and average power with multiple runs of the inference procedure (across tasks).

Remark 4.3. *The* i.i.d. *assumption is formally violated in this experiment, where for each subject we analyze two sample image that are not independent. Yet, this remains a short-range correlation structure, and is thus not a strong challenge to the* i.i.d. *assumption.*
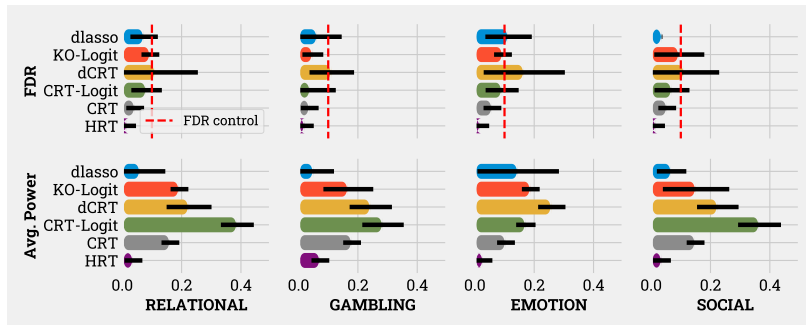

Figure 4: **FDR/Average Power of 50 runs of semi-realistic experiments on four tasks of Human Connectome Project dataset.** Parameters: $n = 800$ (taken from 400 subjects), SNR = 2.0. Methods (clustering versions): Debiased Lasso (`cdlasso`), model-X Knockoff (`cKO-logit`), original CRT (`CRT`), HRT (`HRT`), dCRT (`dCRT`), and our version of CRT (dark green line – `CRT-logit`).

Results Results in Figure 4 show that CRT-logit achieves a better recovery compared to KO or original CRT/dCRT/HRT, which results in higher statistical power. This gain comes with a good control of the FDR under desired level $\alpha = 0.1$. On a related note, the only analysis where dCRT makes more discoveries than CRT-Logit is in `emotion` task, but at the cost of failing to control FDR at nominal level.

4.4. Application: genome-wide association study with Human Brain Cancer Dataset

Description The last in our benchmark is a Genome-wide Association Study (GWAS) on the The Cancer Genome Atlas (TCGA) dataset (Weinstein et al., 2013; Vasaikar et al., 2018). We choose to analyze the Glioma cohort, which consists of $n = 1026$ patients across a wide age range, diagnosed with this type of brain tumor, with a total of $p = 24776$ genes in the data matrix, recorded as copy number variations (CNVs) at the gene level in log ratio format. As with the brain-imaging inference in Section 4.3, we use clustering to reduce the dimension to $C = 1000$ clusters. For the response, a long-term survivor (LTS) is defined as a patient who survived more than five years after diagnosis and would be labeled $y = 0$, and any patient who died within five years would be a short-term survivor (STS), labeled $y = 1$. The objective is to identify significant genes that contribute to classification of the LTS/STS status. Similar to the Human Connectome Project dataset, there is no real ground-truth for the TCGA Glioma. However, we have the list of mutations and the frequency of those detected in the diagnosed patients. We therefore select the 1000 most frequent gene mutations that appeared in this list, *i.e.* the ground truth list consists of 1000 genes (variables).

Result The result from Table 1 shows that CRT-logit finds the largest number of genes. Moreover, most of selected genes in this table are detected in the list of mutated genes found on recorded patients. Some genes are detected by all the benchmarked methods, most prominently `SPEN`, which is found on over 10 % of patients in the cohort.

List of detected genes associated with Glioma Cancer from the TCGA dataset
$n = 1026, p = 24776$. *Empty line (—) signifies no detection. Methods listed in the table are the clustering version. Commonly detected genes between methods are put in bold text. Most detected genes are listed in the mutant list database that can be found in the recorded patients (*Vasaikar et al.*, 2018).*

Methods	Detected Genes
dLasso	—
KO	**ABCC10**, **ANK3**, CDH23, PTEN, **SPEN**, **SVIL**, ZMIZ1
dCRT	**ANK3**, **ANKRD30A**, CDH23, PTEN, RET, **SPEN**,ZMIZ1
CRT-logit	**ABCC10**, **ANKRD30A**, BCOR, EPHA3, PPL, SPAG17, **SPEN**, **SVIL**, USP9X
Original CRT	**ABCC10**, BCOR, EPHA3, **SPEN**, **SVIL**
HRT	**ABCC10**, **SPEN**

Furthermore, this gene is known to be associated not only with brain cancer, but also with other types of cancer in The Human Protein Atlas project (Légaré et al., 2015). Note that, in the absence of a ground-truth, this does not guarantee all genes found are associated with glioma, but this experiment demonstrates the capability of CRT-logit in GWAS studies.

4.5. Average runtime of benchmarked methods

Average runtime of benchmarked methods for one simulation (in seconds). *Standard error is reported in parentheses.*

Methods	Simulated (Sec. 4.2)	HCP-semi-real (Sec. 4.3)
Debiased Lasso (Zhang and Zhang, 2014; van de Geer et al., 2014; Javanmard and Montanari, 2014)	61.83 (5.2)	154.27 (8.79)
Knockoff Filter (Barber and Candès, 2015; Candès et al., 2018)	1.62 (0.02)	8.12 (0.62)
CRT (500 samplings) (Candès et al., 2018)	2312.91 (38.21)	7069.96 (109.09)
HRT (5000 samplings) (Tansey et al., 2022)	14.84 (2.01)	52.17 (4.11)
dCRT[screening=True] (Liu et al., 2021)	16.83 (1.89)	65.18 (3.91)
dCRT[screening=False] (Liu et al., 2021)	370.12 (8.18)	962.40 (20.63)
CRT-logit[screening=True] **(this work)**	**14.16 (0.35)**	**61.26 (3.55)**
CRT-logit[screening=False] **(this work)**	**367.91 (4.11)**	**983.78 (17.26)**

Besides statistical performance, it is equally important to assess the computational cost of inference procedures. The average runtime in Table 2 from the two experiments shows that the original CRT is not suitable for large-scale inference: it is over 2000 times slower than the fastest method (Knockoff Filter), and over 150 times slower than dCRT/CRT-logit. The empirical runtime also confirms the effectiveness of the screening step before doing distillation/decorrelation of the test-statistics: the step makes CRT-logit and dCRT 20 times faster than without. On a related note, although in theory Debiased Lasso, dCRT and CRT-logit (both without screening) share the same runtime complexity, the latter two are slower due to the use of cross-validation to estimate the sparsity hyperparameter λ and λ_{dx} (detailed in Section 3).

5. Discussion

We proposed an adaptation of the Conditional Randomization Test (CRT) for sparse logistic regression in the high-dimensional regime. A major improvement of CRT-logit, our proposed algorithm, compared to original CRT, comes from the decorrelation of

test statistics to make their distribution closer to standard normal. Indeed, results from synthetic experiments in Figure 6 show that in high-dimension (when $0.5 \leq n/p \leq 1.0$), the empirical null distribution of CRT-logit's test statistic T^{decorr} is much more similar to a standard normal compared to the original CRT test statistic. Moreover, empirical benchmarks in Section 4 demonstrate that CRT-logit performs better than related statistical inference methods, such as the Debiased Lasso or Model-X Knockoffs. In particular, CRT-logit is the most powerful method in our synthetic experiment with high-dimensional datasets in Section 4.2, while still keeping FDR controlled under predefined level $\alpha = 0.1$.

We note that there exists some limitations to CRT-logit. The computational cost of CRT-logit, while lower than vanilla CRT, is still larger than alternative methods such as Knockoff Filter and Holdout Randomization Test. Moreover, tuning the ℓ_1-regularization λ_{dx} parameter by cross-validation, as is often done, can further increase the computational cost of CRT-logit (and dCRT).

Despite this, our empirical benchmarks on both simulated and real data show real promises of CRT-logit. Henceforth, we believe CRT-logit is competitive for practical settings that involve structured data, such as brain-imaging and genomics applications.

Acknowledgements

References

Bach, F. (2010). Self-concordant analysis for logistic regression. *Electronic Journal of Statistics*, 4(none).

Bach, F., Jenatton, R., Mairal, J., Obozinski, G., et al. (2012). Optimization with sparsity-inducing penalties. *Foundations and Trends® in Machine Learning*, 4(1):1–106.

Barber, R. F. and Candès, E. J. (2015). Controlling the false discovery rate via knockoffs. *The Annals of Statistics*, 43(5):2055–2085. arXiv: 1404.5609.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300.

Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Ann. Statist.*, 29(4):1165–1188.

Berrett, T. B., Wang, Y., Barber, R. F., and Samworth, R. J. (2020). The conditional permutation test for independence while controlling for confounders. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(1):175–197. _eprint: https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/rssb.12340.

Bzdok, D., Eickenberg, M., Grisel, O., Thirion, B., and Varoquaux, G. (2015). Semi-Supervised Factored Logistic Regression for High-Dimensional Neuroimaging Data. In Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

Candès, E., Fan, Y., Janson, L., and Lv, J. (2018). Panning for gold: 'model-x' knockoffs for high dimensional controlled variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3):551–577.

Chevalier, J.-A., Nguyen, T.-B., Salmon, J., Varoquaux, G., and Thirion, B. (2021). Decoding with confidence: Statistical control on decoder maps. *NeuroImage*, 234:117921.

Cox, D. R. and Hinkley, D. V. (1979). *Theoretical statistics*. CRC Press.

Fu, W. and Knight, K. (2000). Asymptotics for lasso-type estimators. *The Annals of statistics*, 28(5):1356–1378.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The elements of statistical learning: data mining, inference and prediction.* OCLC: 1058138445.

Javanmard, A. and Montanari, A. (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *The Journal of Machine Learning Research*, 15(1):2869–2909.

Légaré, S., Cavallone, L., Mamo, A., Chabot, C., Sirois, I., Magliocco, A., Klimowicz, A., Tonin, P. N., Buchanan, M., Keilty, D., et al. (2015). The estrogen receptor cofactor spen functions as a tumor suppressor and candidate biomarker of drug responsiveness in hormone-dependent breast cancers. *Cancer research*, 75(20):4351–4363.

Lindquist, M. A. (2008). The Statistical Analysis of fMRI Data. *Statistical Science*, 23(4).

Liu, M., Katsevich, E., Janson, L., and Ramdas, A. (2021). Fast and powerful conditional randomization testing via distillation. *Biometrika*.

Nguyen, T.-B., Chevalier, J.-A., and Thirion, B. (2019). Ecko: Ensemble of clustered knockoffs for robust multivariate inference on fmri data. In *International Conference on Information Processing in Medical Imaging*, pages 454–466. Springer.

Ning, Y. and Liu, H. (2017). A general theory of hypothesis tests and confidence regions for sparse high dimensional models. *Annals of Statistics*, 45(1):158–195. Publisher: Institute of Mathematical Statistics.

Ostrovskii, D. M. and Bach, F. (2021). Finite-sample analysis of m-estimators using self-concordance. *Electronic Journal of Statistics*, 15(1):326–391.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830.

Rao, C. R. (1948). Large sample tests of statistical hypotheses concerning several parameters with applications to problems of estimation. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 44, pages 50–57. Cambridge University Press.

Sur, P. and Candès, E. J. (2019). A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116(29):14516–14525. Publisher: Proceedings of the National Academy of Sciences.

Tansey, W., Veitch, V., Zhang, H., Rabadan, R., and Blei, D. M. (2022). The Holdout Randomization Test for Feature Selection in Black Box Models. *Journal of Computational and Graphical Statistics*, 31(1):151–162. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/10618600.2021.1923520.

Thirion, B., Varoquaux, G., Dohmatob, E., and Poline, J.-B. (2014). Which fMRI clustering gives good brain parcellations? *Frontiers in Neuroscience*, 8.

van de Geer, S., Bühlmann, P., Ritov, Y., and Dezeure, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics*, 42(3):1166–1202. arXiv: 1303.0518.

Varoquaux, G., Gramfort, A., and Thirion, B. (2012). Small-sample brain mapping: sparse recovery on spatially correlated designs with randomization and clustering. *Proceedings of the 29th International Conference on Machine Learning, Edinburgh, Scotland, UK, 2012*, page 8.

Vasaikar, S. V., Straub, P., Wang, J., and Zhang, B. (2018). LinkedOmics: analyzing multi-omics data within and across 32 cancer types. *Nucleic Acids Research*, 46(D1):D956–D963.

Weinstein, J. N., Collisson, E. A., Mills, G. B., Shaw, K. R. M., Ozenberger, B. A., Ellrott, K., Shmulevich, I., Sander, C., and Stuart, J. M. (2013). The cancer genome atlas pan-cancer analysis project. *Nature genetics*, 45(10):1113–1120.

Yadlowsky, S., Yun, T., McLean, C. Y., and D' Amour, A. (2021). SLOE: A Faster Method for Statistical Inference in High-Dimensional Logistic Regression. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P. S., and Vaughan, J. W., editors, *Advances in Neural Information Processing Systems*, volume 34, pages 29517–29528. Curran Associates, Inc.

Zhang, C.-H. and Zhang, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):217–242. _eprint: https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/rssb.12026.

Zhao, Q., Sur, P., and Candès, E. J. (2022). The asymptotic distribution of the MLE in high-dimensional logistic models: Arbitrary covariance. *Bernoulli*, 28(3):1835 – 1861. Publisher: Bernoulli Society for Mathematical Statistics and Probability.

Appendix A: Proofs of theoretical results in Section 3

We first present some technical lemmas that are useful for the proof of the main theorem. From now on, let $\precsim$ and $\succsim$ denote inequalities with a hidden constant factor, *i.e.* $x \precsim y$ means that with high probability, there exists an absolute constant $C > 0$ such that $x \leq Cy$, and vice versa. As mentioned in the main text, in what follows, without writing it explicitly, we consider $p = p(n)$.

Lemma A.1 (Lemma E.1, Ning and Liu (2017))**.** *Assume Assumption* **??***, under the logistic model, we have*

$$\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^0\|_1 \precsim s^* \sqrt{\frac{\log p}{n}} \quad and \quad \|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^0\|_2 \precsim \sqrt{\frac{s^* \log p}{n}},$$

where $s^* = \|\boldsymbol{\beta}^0\|_0$*. In addition, we also have*

$$\frac{1}{n} \sum_{i=1}^{n} g''(\mathbf{X}_{i,*}\boldsymbol{\beta}^0)[\mathbf{X}_{i,-j}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^0)]^2 \precsim \frac{s^* \log(p)}{n} ,$$

where $g(x) = 1/(1 + \exp(x))$ *is the sigmoid function.*

Lemma A.2 (Lemma E.2 Ning and Liu (2017), concentration of the gradient and Hessian of the logistic loss function)**.** *Assume Assumption* **??** *holds, under logistic model, we have, with* $\mathbf{v}^* \overset{def.}{=} (1, -\mathbf{w}^{0,j}) \in \mathbb{R}^p$*,*

$$\|\nabla\ell(\boldsymbol{\beta}^0)\|_\infty \precsim \sqrt{n^{-1}\log p}, \quad and$$
$$\|\mathbf{v}^{*\top}\nabla^2\ell(\boldsymbol{\beta}^0) - \mathbb{E}_{\boldsymbol{\beta}^0}[\mathbf{v}^{*\top}\nabla^2\ell(\boldsymbol{\beta}^0)]\|_\infty \precsim \sqrt{n^{-1}\log p}.$$

Lemma A.3 (Lemma E.3, Ning and Liu (2017))**.** *Assume Assumption* **??** *holds, under logistic model, we have*

$$\|\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} - \mathbf{w}^{0,j}\|_1 \precsim (s' \vee s^*)\sqrt{\frac{\log p}{n}},$$

where $s^* = \|\boldsymbol{\beta}^0\|_0$ *and* $s' = \|\mathbf{w}^{0,j}\|_0$*. In addition, we also have*

$$\frac{1}{n} \sum_{i=1}^{n} g''(\mathbf{X}_{i,*}\hat{\boldsymbol{\beta}})[\mathbf{X}_{i,-j}(\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} - \mathbf{w}^{0,j})]^2 \precsim \frac{(s' \vee s^*)\log(p)}{n}.$$

Lemma A.4 (Lemma E.4, Ning and Liu (2017), local smoothness conditions on the loss function)**.** *Let* $\hat{\boldsymbol{\beta}}^{null} = (0, \hat{\boldsymbol{\beta}}_{-j}) \in \mathbb{R}^p$*, where* $\hat{\boldsymbol{\beta}}$ *is an estimator of* $\boldsymbol{\beta}^0$*. It holds that*

$$\mathbf{v}^{*\top}[\nabla\ell(\boldsymbol{\beta}) - \nabla\ell(\boldsymbol{\beta}^0) - \nabla^2\ell(\boldsymbol{\beta}^0)(\boldsymbol{\beta} - \boldsymbol{\beta}^0)] \precsim \frac{(s^* \vee s')\log p}{n} ,$$
$$(\hat{\mathbf{v}} - \mathbf{v}^*)^\top[\nabla\ell(\boldsymbol{\beta}) - \nabla\ell(\boldsymbol{\beta}^0)] \precsim \frac{(s^* \vee s')\log p}{n} .$$

for both $\boldsymbol{\beta} = \hat{\boldsymbol{\beta}}^{null}$ *and* $\boldsymbol{\beta} = \hat{\boldsymbol{\beta}}$*.*

Proof of Theorem 3.1. The following proof is an adaptation from Ning and Liu (2017). Notice that our version of the proof is shorter, with specific consideration on sparse logistic regression, and with elaboration on the convergence rate of the decorreleated test score, which is missing from Ning and Liu (2017).

Denote $\hat{\mathbf{v}} \overset{def.}{=} (1, (\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}})^\top)$, then the decorrelated test score can be written in more general from as

$$T_j^{\text{decorr}} = n^{1/2}\hat{\mathbf{I}}_{j|-j}^{-1/2}\left(\nabla_j\ell(\hat{\boldsymbol{\beta}}) - (\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}})^\top\nabla_{\boldsymbol{\beta}_{-j}}\ell(\hat{\boldsymbol{\beta}})\right) = n^{1/2}\hat{\mathbf{I}}_{j|-j}^{-1/2}\hat{\mathbf{v}}^\top\nabla\ell(\hat{\boldsymbol{\beta}}) . \quad (14)$$

Moreover, denote $\hat{\boldsymbol{\beta}}^{\text{null}} \overset{def.}{=} (0, \hat{\boldsymbol{\beta}}_{-j})$ and $\mathbf{v}^* \overset{def.}{=} (1, \mathbf{w}^{0,j})$, then we have, under the null hypothesis,

$$n^{1/2}|\hat{\mathbf{v}}^\top\nabla(\hat{\boldsymbol{\beta}}^{\text{null}}) - \mathbf{v}^{*\top}\nabla\ell(\boldsymbol{\beta}^0)| \leq \underbrace{n^{1/2}|\mathbf{v}^{*\top}\{\nabla\ell(\boldsymbol{\beta}^0) - \nabla\ell(\hat{\boldsymbol{\beta}}^{\text{null}})\}|}_{A_1} + \underbrace{n^{1/2}|(\hat{\mathbf{v}} - \mathbf{v}^*)^\top\nabla\ell(\hat{\boldsymbol{\beta}}^{\text{null}})|}_{A_2}$$

where we use triangle inequality with the last step. By Taylor expansion, and from Lemma A.4, we have

$$
\begin{aligned}
A_1 &\leq n^{1/2} |\mathbf{v}^{*\top} \nabla^2 \ell(\boldsymbol{\beta}^0)(\hat{\boldsymbol{\beta}}^{\text{null}} - \boldsymbol{\beta}^0)| \\
&\leq n^{1/2} \|\hat{\boldsymbol{\beta}}^{\text{null}} - \boldsymbol{\beta}^0\|_1 \|\mathbf{v}^{*\top} \nabla^2 \ell(\boldsymbol{\beta}^0)\|_\infty \\
&\precsim \frac{s^* \log(p)}{\sqrt{n}}
\end{aligned}
$$

where the second inequality is by Holdër inequality, and the last inequality us due to Lemma A.1 and A.2. Similarly, we can bound A_2, by using Lemma A.3 and Lemma A.4

$$
\begin{aligned}
A_2 &\leq n^{1/2} |(\hat{\mathbf{v}} - \mathbf{v}^*)^\top \nabla \ell(\boldsymbol{\beta}^0)| \\
&\leq n^{1/2} \|\hat{\mathbf{v}} - \mathbf{v}^*\|_1 \|\nabla \ell(\boldsymbol{\beta}^0)\|_\infty \precsim \frac{(s^* \vee s') \log p}{n}
\end{aligned}
$$

This implies that,

$$
n^{1/2} |\hat{\mathbf{v}}^\top \nabla(\hat{\boldsymbol{\beta}}^{\text{null}}) - \mathbf{v}^{*\top} \nabla \ell(\boldsymbol{\beta}^0)| \precsim n^{-1/2}(s^* \vee s') \log(p) . \tag{15}
$$

The remaining part of the proof is to bound $\hat{\mathbf{I}}_{j|-j} - \mathbf{I}_{j|-j}$, where, by definition

$$
\mathbf{I}_{j|-j} = \mathbb{E} \left\{ g''(\mathbf{X}_{i,*} \boldsymbol{\beta}^0) \left[\mathbf{X}_{i,j} - \mathbf{X}_{i,-j} \mathbf{w}^{0,j} \right] \mathbf{X}_{i,j} \right\}
$$

Evaluating the difference between $\hat{\mathbf{I}}_{j|-j}$ and $\mathbf{I}_{j|-j}$ gives

$$
\begin{aligned}
&\hat{\mathbf{I}}_{j|-j} \text{ - } \mathbf{I}_{j|-j} \\
&= \frac{1}{n} \sum_{i=1}^n g''(\mathbf{X}_{i,*} \hat{\boldsymbol{\beta}}) \left[\mathbf{X}_{i,j} - \mathbf{X}_{i,-j} \hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} \right] \mathbf{X}_{i,j} - \mathbb{E} \left\{ g''(\mathbf{X}_{i,*} \boldsymbol{\beta}^0) \left[\mathbf{X}_{i,j} - \mathbf{X}_{i,-j} \mathbf{w}^{0,j} \right] \mathbf{X}_{i,j} \right\} \\
&= \left(\frac{1}{n} \sum_{i=1}^n g''(\mathbf{X}_{i,*} \hat{\boldsymbol{\beta}}) \mathbf{X}_{i,j}^2 - \mathbb{E} \left\{ g''(\mathbf{X}_{i,*} \boldsymbol{\beta}^0) \mathbf{X}_{i,j}^2 \right\} \right) \\
&\quad + \left(\frac{1}{n} \sum_{i=1}^n g''(\mathbf{X}_{i,*} \hat{\boldsymbol{\beta}}) \mathbf{X}_{i,-j} \hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} \mathbf{X}_{i,j} - \mathbb{E} \left\{ g''(\mathbf{X}_{i,*} \boldsymbol{\beta}^0) \mathbf{X}_{i,-j} \mathbf{w}^{0,j} \mathbf{X}_{i,j} \right\} \right) \\
&\leq \underbrace{\left(\frac{1}{n} \sum_{i=1}^n g''(\mathbf{X}_{i,*} \hat{\boldsymbol{\beta}}) \mathbf{X}_{i,j}^2 - \mathbb{E} \left\{ g''(\mathbf{X}_{i,*} \boldsymbol{\beta}^0) \mathbf{X}_{i,j}^2 \right\} \right)}_{C} + \underbrace{\left| \frac{1}{n} \sum_{i=1}^n g''(\mathbf{X}_{i,*} \hat{\boldsymbol{\beta}}) \mathbf{X}_{i,-j} (\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} - \mathbf{w}^{0,j}) \mathbf{X}_{i,j} \right|}_{B_1} \\
&\quad + \underbrace{\left| \frac{1}{n} \sum_{i=1}^n [g''(\mathbf{X}_{i,*} \boldsymbol{\beta}^0) - g''(\mathbf{X}_{i,*} \hat{\boldsymbol{\beta}})] \mathbf{X}_{i,-j} \mathbf{w}^{0,j} \mathbf{X}_{i,j} \right|}_{B_2} + \\
&\quad \underbrace{\left| \frac{1}{n} \sum_{i=1}^n g''(\mathbf{X}_{i,*} \boldsymbol{\beta}^0) \mathbf{X}_{i,-j} \mathbf{w}^{0,j} \mathbf{X}_{i,j} - \mathbb{E} \left\{ g''(\mathbf{X}_{i,*} \boldsymbol{\beta}^0) \mathbf{X}_{i,-j} \mathbf{w}^{0,j} \mathbf{X}_{i,j} \right\} \right|}_{B_3},
\end{aligned}
$$

where the last step follows triangle inequality.

We have, by Cauchy-Schwartz inequality, by Lemma A.3, and by the fact that $g''(x) \in (0,1)$ for every $x \in \mathbb{R}$; and $\mathbf{X}_{i,-j}, \mathbf{X}_{i,j}$ is sub-exponential by Assumption 3.1:

$$
\begin{aligned}
B_1 &\leq \sqrt{\left(\frac{1}{n} \sum_{i=1}^n g''(\mathbf{X}_{i,*}^\top \hat{\boldsymbol{\beta}}) ((\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} - \mathbf{w}^{0,j})^\top \mathbf{X}_{i,-j})^2 \right) \left(\frac{1}{n} \sum_{i=1}^n g''(\mathbf{X}_{i,*}^\top \hat{\boldsymbol{\beta}}) \mathbf{X}_{i,j}^2 \right)} \\
&\precsim \sqrt{\frac{(s^* \vee s') \log(p)}{n}}.
\end{aligned}
$$

Similarly, to bound B_2, we have, again by Cauchy-Schwartz inequality,

$$B_2 \leq \sqrt{\frac{1}{n}\sum_{i=1}^{n}[g''(\mathbf{X}_{i,*}\boldsymbol{\beta}^0) - g''(\mathbf{X}_{i,*}\hat{\boldsymbol{\beta}})]^2 (\mathbf{X}_{i,-j}\mathbf{w}^{0,j} \ \mathbf{X}_{i,j})^2}$$

$$\leq \sqrt{\frac{1}{n}\sum_{i=1}^{n}[g''(\mathbf{X}_{i,*}\boldsymbol{\beta}^0)\mathbf{X}_{i,*}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^0)]^2 (\mathbf{X}_{i,-j}\mathbf{w}^{0,j} \ \mathbf{X}_{i,j})^2} \ ,$$

where the second inequality comes from using the self-concordance property of the sigmoid function (discussed at length in Bach (2010) and extended further in Ostrovskii and Bach (2021)), that is, $|g''(t_1) - g''(t)| \leq |t_1 - t|g''(t)|$ for a fixed constant t, and for every $t_1 \in \mathbb{R}$ such that t_1 converges to t, with $t_1 = \hat{\boldsymbol{\beta}}$, and $t = \boldsymbol{\beta}^0$. By Assumption 3.1-A3 that $\mathbf{X}_{i,j}$ is sub-exponential, applying Bernstein inequality leads to

$$B_2 \precsim \sqrt{\frac{s^* \log p}{n}}.$$

To bound B_3, by direct application of Hoeffding inequality, we have $B_3 \precsim \sqrt{\frac{(s^* \vee s') \log p}{n}}$. This implies

$$|\hat{\mathbf{I}}_{j|-j} - \mathbf{I}_{j|-j}| \precsim \sqrt{\frac{(s^* \vee s') \log p}{n}} \ . \tag{16}$$

Putting Equation (15) and (16) together, we have, under null hypothesis,

$$T_j^{\text{decorr}} \xrightarrow{\mathcal{D}} n^{1/2} \ \mathbf{I}_{j|-j}^{-1/2}\mathbf{v}^{*\top}\nabla\ell(\boldsymbol{\beta}^0) \overset{\text{def.}}{=} T_j^* \ ,$$

with convergence rate $\mathcal{O}(n^{-1/2})$. Finally, by noting that we can decompose $\nabla\ell(\boldsymbol{\beta}^0) = \frac{1}{n}\sum_{i=1}^{n}\nabla\ell_i(\boldsymbol{\beta}^0)$, and each $\ell_i(\boldsymbol{\beta}^0)$ has bounded first, second, and third moment, a direct application of Berry-Esseen theorem give convergence in distribution of T_j^* to a standard normal law, with rate $\mathcal{O}(n^{-1/2})$.

We also arrive at the second conclusion of Theorem 3.1 by noting that it is a straightforward by-product of the result on normality of the distribution of decorrelated test score under null hypothesis, based on the formula for the p-values of CRT-logit algorithm. $\qquad\square$

Proof of Theorem 3.2. The proof of this theorem is a straightforward adaptation from Benjamini and Yekutieli (2001). For shorter notation, we denote $\widehat{\mathcal{S}} \overset{\text{def.}}{=} \widehat{\mathcal{S}}_{\text{BY-CRT}}$ and $\widehat{k} \overset{\text{def.}}{=} \widehat{k}_{BY}$. If we denote $\bar{\alpha} \overset{\text{def.}}{=} \frac{\alpha}{p\sum_{i=1}^{p}1/i} \in (0,1)$, then step 1 in the procedure defined in Definition B.2 is equivalent with finding $\widehat{k}$ such that

$$\widehat{k} = \max\left\{k \in [p] \mid \hat{p}_{(k)} \leq k\bar{\alpha}\right\}. \tag{17}$$

For every $i, j, k \in [p]$, let us define

$$p_{i,j,k} = \begin{cases} \mathbb{P}\left(\widehat{p}_i \in ((j-1)\bar{\alpha}, j\bar{\alpha}] \ , i \in \widehat{S} \text{ and } |\widehat{S}| = k\right) & \text{if } j \geq 2 \\ \mathbb{P}\left(\widehat{p}_i \in [0, \bar{\alpha}] \ , i \in \widehat{S} \text{ and } |\widehat{S}| = k\right) & \text{if } j = 1. \end{cases} \tag{18}$$

Then, since $i \in \widehat{\mathcal{S}}$ and $|\widehat{\mathcal{S}}| = k$ implies that $\widehat{p}_i \leq \widehat{p}_{\widehat{k}} \leq \widehat{k}\bar{\alpha} = k\bar{\alpha}$, we have

$$\frac{|\widehat{S} \cap \mathcal{S}^c|}{|\widehat{S}| \vee 1} = \sum_{k=1}^{p}\mathbb{1}_{|\widehat{S}|=k}\frac{\sum_{i \in \mathcal{S}^c}\mathbb{1}_{i \in \widehat{S}}}{k} = \sum_{i \in \mathcal{S}^c}\sum_{k=1}^{p}\frac{1}{k}\mathbb{1}_{|\widehat{S}|=k \text{ and } i \in \widehat{S}}$$

$$= \sum_{i \in \mathcal{S}^c}\sum_{k=1}^{p}\frac{1}{k}\mathbb{1}_{|\widehat{S}|=k \text{ and } i \in \widehat{S} \text{ and } 0 \leq \hat{p}_i \leq k\bar{\alpha}} \ .$$

Taking an expectation and writing that

$$\mathbb{1}_{0 \le \widehat{p}_i \le k\bar{\alpha}} = \mathbb{1}_{\widehat{p}_i \in [0,\bar{\alpha}]} + \sum_{j=2}^{k} \mathbb{1}_{\widehat{p}_i \in ((j-1)\bar{\alpha}, j\bar{\alpha}]} \ ,$$

we get

$$\mathbb{E}\left[\frac{|\widehat{\mathcal{S}} \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}| \vee 1} \right] = \sum_{i \in \mathcal{S}^c} \sum_{k=1}^{p} \frac{1}{k} \sum_{j=1}^{k} p_{i,j,k} = \sum_{i \in \mathcal{S}^c} \sum_{j=1}^{p} \sum_{k=j}^{p} \frac{1}{k} p_{i,j,k}$$

$$\le \sum_{i \in \mathcal{S}^c} \sum_{j=1}^{p} \sum_{k=j}^{p} \frac{1}{j} p_{i,j,k} = \underbrace{\sum_{j=1}^{p} \frac{1}{j} \sum_{i \in \mathcal{S}^c} \sum_{k=j}^{p} p_{i,j,k}}_{A} \ .$$

Denote $F(j) \overset{\text{def.}}{=} \sum_{i \in \mathcal{S}^c} \sum_{j'=1}^{j} \sum_{k=1}^{p} p_{i,j',k}$ for all $j \in \{1,\dots,p\}$, and remark that $p_{i,j',k} = 0$ if $j' > k$, by definition of $\widehat{\mathcal{S}}_{\text{BY-CRT}}$. We then have

$$A = F(1) + \sum_{j=2}^{p} \frac{1}{j} \left[F(j) - F(j-1) \right] = \sum_{j=1}^{p-1} \left(\frac{1}{j} - \frac{1}{j+1} \right) F(j) + \frac{F(p)}{p} \ .$$

This leads to

$$\mathbb{E}\left[\frac{|\widehat{\mathcal{S}} \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}| \vee 1} \right] \le \sum_{j=1}^{p-1} \left(\frac{1}{j} - \frac{1}{j+1} \right) F(j) + \frac{F(p)}{p} \tag{19}$$

By the definition of $p_{i,j,k}$ in Eq. (18), we have

$$F(j) = \sum_{i \in \mathcal{S}^c} \mathbb{P}(\widehat{p}_i \le j\bar{\alpha} \text{ and } i \in \widehat{\mathcal{S}}) \le \sum_{i \in \mathcal{S}^c} \mathbb{P}(\widehat{p}_i \le j\bar{\alpha}).$$

Therefore

$$\mathbb{E}\left[\frac{|\widehat{\mathcal{S}} \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}| \vee 1} \right] \le \sum_{i \in \mathcal{S}^c} \sum_{j=1}^{p-1} \frac{\mathbb{P}(\widehat{p}_i \le j\bar{\alpha})}{j(j+1)} + \sum_{i \in \mathcal{S}^c} \frac{\mathbb{P}(\widehat{p}_i \le p\bar{\alpha})}{p}$$

Taking the limit where $n \to \infty$ and p fixed, we have, using the result in Theorem 3.1,

$$\limsup_{n \to \infty} \mathbb{E}\left[\frac{|\widehat{\mathcal{S}} \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}| \vee 1} \right] \le \sum_{i \in \mathcal{S}^c} \left(\sum_{j=1}^{p-1} \frac{1}{j+1} + 1 \right) \bar{\alpha}$$

$$= \left(\sum_{j=1}^{p} \frac{1}{j} \right) |\mathcal{S}^c| \bar{\alpha} \ .$$

We conclude the proof by noting that $\bar{\alpha} \overset{\text{def.}}{=} \dfrac{\alpha}{p \sum_{j=1}^{p} 1/j}$. $\qquad\square$

Appendix B: Controlling False Discovery Rate Procedures

Definition B.1 (Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995))**.** *Let $\alpha \in (0,1)$ be the predefined FDR control level. Let $\hat{p}_1, \dots, \hat{p}_m$ be output p-values from inference algorithm,* e.g. *Algorithm 1. We reorder them ascendingly, denoted by $\hat{p}_{(1)} \le \hat{p}_{(2)} \le \dots \le \hat{p}_{(p)}$ and $\mathcal{H}_0^{(1)}, \dots, \mathcal{H}_0^{(p)}$, then*

1. *Find $\widehat{k}_{BH}$ such that*

$$\widehat{k}_{BY} = \max\left\{ k \in [p] \mid \hat{p}_{(k)} \le \frac{k\alpha}{p} \right\} .$$

2. *If $\widehat{k}_{BH}$ exists, take $\widehat{\mathcal{S}} = \{ j \in [p] : \hat{p}_{(j)} \le \hat{p}_{\widehat{k}_{BH}} \}$. Otherwise $\widehat{\mathcal{S}} = \emptyset$.*

Definition B.2 (Benjamini-Yekutieli procedure (Benjamini and Yekutieli, 2001))**.** *Let* $\alpha \in (0,1)$ *be the predefined FDR control level. Let* $\hat{p}_1, \ldots, \hat{p}_m$ *be output p-values from Algorithm 1. We reorder them ascendingly, denoted by* $\hat{p}_{(1)} \leq \hat{p}_{(2)} \leq \cdots \leq \hat{p}_{(p)}$ *and* $\mathcal{H}_0^{(1)}, \ldots, \mathcal{H}_0^{(p)}$, *then*

1. Find $\widehat{k}_{BY}$ *such that*

$$\widehat{k}_{BY} = \max\left\{ k \in [p] \mid \hat{p}_{(k)} \leq \frac{k\alpha}{p \sum_{i=1}^{p} 1/i} \right\}.$$

2. If $\widehat{k}_{BY}$ *exists, take* $\widehat{S} = \{j \in [p] : \hat{p}_{(j)} \leq \hat{p}_{\widehat{k}_{BY}}\}$. *Otherwise* $\widehat{S} = \emptyset$.

Appendix C: Setting the ℓ_1-Regularization Parameter of the $\mathbf{X}_{*,j}$-distillation

A core issue is the dependency of the statistical power and FDR of CRT-logit on the ℓ_1- regularization parameter λ_{dx} when doing Lasso distillation on x_j in Eq. (10). One might choose the heuristic value $\lambda_{\text{univ}} = \sqrt{n^{-1} \log p}$ with theoretical validity, as suggested in Ning and Liu (2017); van de Geer et al. (2014). However, experimental results in Fig. 5 show that at $\lambda_{dx} = \lambda_{\text{univ}}$ (or $\log_{10} \lambda/\lambda_{\text{univ}} = 0.0$ with the labeling of the figure), we do not have the best possible FDR/Power with CRT-logit inference. For this experiment, we average the inference results of 100 simulations (with similar setting in Section 4.1) for different values of n and λ_{dx}, with p fixed. There is a clear phase transition in both FDR and average power when the regularization parameter λ_{dx} increases. In other words, we have found empirically that both FDR and power of the method are sensitive to the ℓ_1-regularization parameter. Preferably, one wants to return a high statistical power while controlling FDR under predefined level. Hence, it is necessary to choose λ_{dx} wisely. In a more practical scenario, we advise to use cross-validation for $\mathbf{X}_{*,j}$-distillation operator, as defined by Eq. (10). This means we would have to find p different values of λ_{dx} with cross-validation, and we reemphasize the importance of the screening step to reduce the number of computations.


Figure 5: **FDR/Average Power of 100 runs of simulations while varying the number of samples and ℓ_1 regularization parameter and fixing the number of variables**. Note: λ_{dx} is scaled with the factor $\lambda_{\text{univ}} = \sqrt{\log(p)/n}$, *e.g.* the first value for regularization grid is $\lambda_{dx} = 10^{-2}\lambda_{\text{univ}}$. Default parameter (similar settings in Section 4.1): $p = 400$, SNR=3.0 (signal-to-noise ration), $\rho = 0.5$ (feature correlation), $\kappa = 0.05$ (sparsity). FDR is controlled at level $\alpha = 0.1$.

Appendix D: Pseudocode for CRT-logit and Related Algorithms

Appendix E: Time complexity of Related Methods

We present the time complexity of benchmarked methods in Table 3.

Algorithm 2: Conditional Randomization Test (Candès et al., 2018)

---

1 **INPUT** dataset $(\mathbf{X}, \mathbf{y})$, with $\mathbf{X} \in \mathbb{R}^{n \times p}$, $\mathbf{y} \in \mathbb{R}^n$, number of sampling runs B, test statistic T_j, conditional distribution $P_{j|-j}$ for each $j = 1, \ldots, p$;

2 **OUTPUT** vector of p-values $\{\hat{p}_j\}_{j=1}^p$;

3 **for** $j = 1, 2, \ldots, p$ **do**

4   Compute test statistics T_j for original variable;

5   **for** $b = 1, 2, \ldots, B$ **do**

6   1. Generate $\tilde{\mathbf{X}}_{*,j}^{(b)}$, a knockoff sample from $P_{j|-j}$;

7   2. Compute $\tilde{T}_j^{(b)}$ for knockoff variables;

8   **end**

9   Compute the empirical p-value

$$\widehat{p}_j = \frac{1 + \sum_{b=1}^B \mathbf{1}_{\tilde{T}_j^{(b)} \geq T_j}}{1 + B}$$

10 **end**

Algorithm 3: Lasso-Distillation Conditional Randomization Test (Liu et al., 2021)

---

1 **INPUT** dataset $(\mathbf{X}, \mathbf{y})$, $X \in \mathbb{R}^{n \times p}$, $\mathbf{y} \in \mathbb{R}^n$, test statistic T_j for each $j = 1, \ldots, p$;

2 **OUTPUT** vector of p-values $\{p_j\}_{j=1}^p$;

3 $\hat{\mathcal{S}}^{\text{SCREENING}} = \{j : j \in [p], \hat{\beta}_j^{\text{MLE}} \neq 0\}$ `// Using Eq.` (2)

4 **for** $j \in \hat{\mathcal{S}}^{SCREENING}$ **do**

5   1. Distill information of $\mathbf{X}_{-j}$ to $\mathbf{X}_{*,j}$ and to $\mathbf{y}$ by finding:

- $\hat{\boldsymbol{\beta}}^{d_y,j}(\lambda) \leftarrow$ `solve_sparse_logistic_cv`$(\mathbf{X}_{-j}, \mathbf{y})$ `// Using Eq.` (2)

- $\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}}(\lambda) = \operatorname{argmin}_{\boldsymbol{\beta} \in \mathbb{R}^{p-1}} \frac{1}{2} \|\mathbf{X}_{*,j} - \mathbf{X}_{-j}\boldsymbol{\beta}\|_2^2 + \lambda_{dx} \|\boldsymbol{\beta}\|_1$ `// with` λ_{dx} `set using`
  `cross-validation`

    2. Obtain test statistic:

$$T_j = \sqrt{n} \frac{\langle \mathbf{y} - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d_y,j}, \mathbf{X}_{*,j} - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} \rangle}{\left\| \mathbf{y} - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d_y,j} \right\|_2 \left\| \mathbf{X}_{*,j} - \mathbf{X}_{-j}\hat{\boldsymbol{\beta}}^{d\mathbf{x}_{*,j}} \right\|_2}$$

    3. Compute (two-sided) p-value

$$\widehat{p}_j = 2\left[1 - \Phi\left(T_j\right)\right]$$

6 **end**

Appendix F: Additional Details on Experiments in Section 4

F.1. Preprocessing of the brain-imaging dataset

The Human Connectome Project dataset (HCP) is a collection of brain imaging data on healthy young adult subjects with age ranging from 22 to 35. The participants performed different tasks while being scanned by a magnetic resonance imaging (MRI) device to record blood oxygenation level dependent (BOLD) signals of the brain. The aim of this analysis is to investigate which areas of the brain can predict cognitive activity across participants, while taking into account the information from other brain regions. The brain imaging modalities include, among others, resting-state fMRI (R-fMRI) and task-evoked fMRI (T-fMRI). In this work, we only deal with decoding the task-evoked fMRI dataset. The four classification problems we are working with are as follows.

- Relational: predict whether the participant matches figures or identified feature similarities.
- Gambling: predict whether the participant gains or loses gambles.

Algorithm 4: Holdout Randomization Test ([Tansey et al.](), [2022]())

---

1 **INPUT** dataset $(\mathbf{X}, \mathbf{y})$, with $\mathbf{X} \in \mathbb{R}^{n \times p}$, $\mathbf{y} \in \mathbb{R}^n$, number of sampling runs B, test statistic T_j, conditional distribution $P_{j|-j}$ for each $j = 1, \ldots, p$, empirical risk $L(\cdot)$;

2 **OUTPUT** vector of p-values $\{\hat{p}_j\}_{j=1}^p$;

3 $(\mathbf{X}_{\text{train}}, \mathbf{y}_{\text{train}}), (\mathbf{X}_{\text{test}}, \mathbf{y}_{\text{test}}) \leftarrow \texttt{data\_splitting}(\mathbf{X}, \mathbf{y})$;

4 $\hat{f}_\theta \leftarrow \texttt{model\_fitting}(\mathbf{X}_{\text{train}}, \mathbf{y}_{\text{train}})$;

5 **for** $j = 1, 2, \ldots, p$ **do**

6 $\quad T_j \leftarrow L(\mathbf{X}_{\text{test}}, \mathbf{y}_{\text{test}}, \hat{f}_\theta(\mathbf{X}_{\text{test}}))$;

7 $\quad$ **for** $b = 1, 2, \ldots, B$ **do**

8 $\quad\quad$ 1. Generate $\tilde{\mathbf{X}}_{*,j}^{(b)} \sim P_{j|-j}$;

9 $\quad\quad$ 2. $\tilde{T}_j^{(b)} \leftarrow L(\tilde{\mathbf{X}}_{*,j}^{(b)}, \mathbf{y}_{\text{test}}, \hat{f}_\theta(\tilde{\mathbf{X}}_{*,j}^{(b)}))$;

10 $\quad$ **end**

11 $\quad$ Compute the empirical p-value

$$\hat{p}_j = \frac{1 + \sum_{b=1}^B \mathbf{1}_{\tilde{T}_j^{(b)} \geq T_j}}{1 + B}$$

12 **end**

TABLE 3
Time complexities of related methods with CRT-logit, *where p is the dimension size (number of variables), B is the number of sampling runs, and $\hat{k} \ll p$ the cardinality of the screening set (see Section D for more details).*

Methods	Time (Iteration) Complexity	References
Debiased Lasso	$\mathcal{O}(p^4)$	Zhang and Zhang (2014); van de Geer et al. (2014); Javanmard and Montanari (2014)
Knockoff Filter	$\mathcal{O}(p^3)$	Barber and Candès (2015); Candès et al. (2018)
CRT	$\mathcal{O}(Bp^4)$	Candès et al. (2018)
HRT	$\mathcal{O}(p^3 + Bp^2)$	Tansey et al. (2022)
dCRT (with screening)	$\mathcal{O}(\hat{k}p^3)$	Liu et al. (2021)
CRT-logit (with screening)	$\mathcal{O}(\hat{k}p^3)$	**(this work)**

- Emotion: predict whether the participant watches an angry face or a geometric shape.
- Social: predict whether the participant watches a movie with social behavior or not.

To perform dimension reduction, our goal is to apply a clustering scheme that keeps the spatial structure of the data. This is achieved with data-driven parcellation along with a spatially constrained clustering algorithm, following the conclusions by Varoquaux et al. (2012) and Thirion et al. (2014). The hierarchical clustering scheme that we use recursively merges pair of clusters of features based on a criterion that minimized the within-cluster variance. This algorithm is implemented in `scikit-learn` Pedregosa et al. (2011), a popular package for applied machine learning.

Appendix G: Example of decoding maps in semi-realistic brain-analysis experiment of Section 4.3

(a) Debiased Lasso (`dlasso`)   (b) Knockoffs (`KO`)   (c) CRT (`CRT`), $B = 500$

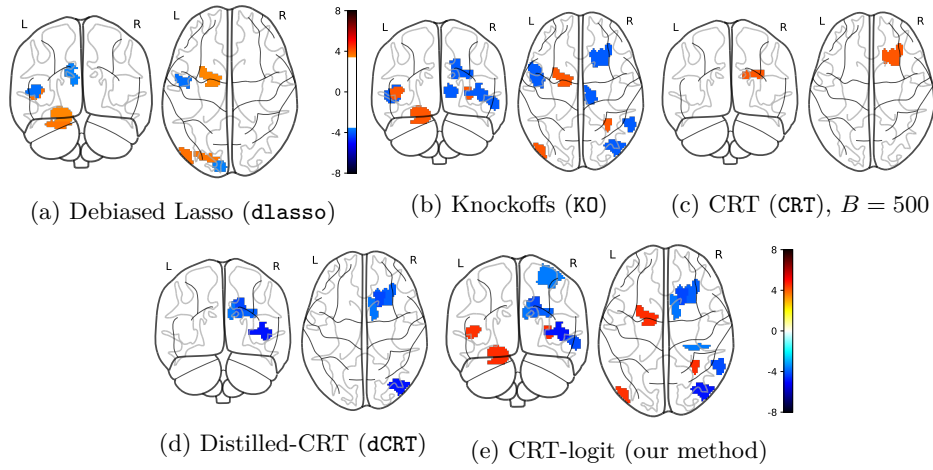(d) Distilled-CRT (`dCRT`)   (e) CRT-logit (our method)

Figure 6: **Decoding maps of Relational task in semi-realistic HCP900 experiment, using 400 subjects and dimension reduction to 1000 clusters (*i.e.* one random seed for generating labels y)**. We omit Holdout Randomization Test (`HRT`) as the method does not select any brain region. For `dlasso`, `dCRT` and `CRT-logit`, we plot the test-statistics; for `KO` the sign of selected coefficients, and for `CRT` the $-log_{10}$ of the empirical p-values.