Estimation and Testing for Multiple Regulation of Multivariate Mixed Outcomes

Denis Agniel,¹* Katherine P. Liao,² and Tianxi Cai³

¹Department of Biomedical Informatics, Harvard Medical School, Boston, Massachusetts, U.S.A. 02115 ²Brigham and Women's Hospital, Boston, Massachusetts, U.S.A. 02115

³Department of Biostatistics, Harvard School of Public Health, Boston, Massachusetts, U.S.A. 02115

 $^{*}email:$ denis.agniel@mail.harvard.edu

SUMMARY. Considerable interest has recently been focused on studying multiple phenotypes simultaneously in both epidemiological and genomic studies, either to capture the multidimensionality of complex disorders or to understand shared etiology of related disorders. We seek to identify *multiple regulators* or predictors that are associated with multiple outcomes when these outcomes may be measured on very different scales or composed of a mixture of continuous, binary, and not-fully observed elements. We first propose an estimation technique to put all effects on similar scales, and we induce sparsity on the estimated effects. We provide standard asymptotic results for this estimator and show that resampling can be used to quantify uncertainty in finite samples. We finally provide a multiple testing procedure which can be geared specifically to the types of multiple regulators of interest, and we establish that, under standard regularity conditions, the familywise error rate will approach 0 as sample size diverges. Simulation results indicate that our approach can improve over unregularized methods both in reducing bias in estimation and improving power for testing.

KEY WORDS: Familywise error control; Hierarchical lasso; Multiple phenotypes; Multiple regulation; Multiple testing; Resampling; Semiparametric models; Stepdown testing.

1. Introduction

Considerable recent interest has been focused on studying multiple phenotypes simultaneously in both epidemiological and genomic studies. There are several reasons for such studies to be important. First, a complex disorder is usually associated with multiple correlated phenotypes. Hence, even when the focus of the study is on a single disease, multiple phenotypes might be needed to fully capture the complexity and multidimensionality of the disorder. Second, multiple related disorders might share the same etiology and a joint assessment will enable researchers to identify factors associated with risk of multiple diseases. As an example, recent studies have identified common genes associated with a higher risk of what were previously considered distinct autoimmune diseases (Zhernakova et al., 2009). Similar shared genetic bases have also been suggested for various types of cancers and related psychiatric disorders (Solovieff et al., 2013). Identification of predictors of multiple outcomes, also commonly known as multiple traits in the genetics literature, can improve understanding of disease etiology, genetic regulatory pathways, and treatment. Further complicating matters, the outcome measures may be *diverse*: they may be binary (e.g., presence of disease), continuous (disease activity score), ordinal (severity of disease), not completely observable (perhaps due to a limit of quantification), or any combination thereof.

To address these questions statistically, we seek to assess the association between a vector of predictors $\mathbf{x} = (x_1, \ldots, x_p)^T$ and a vector of outcomes $\mathbf{y} = (y^{(1)}, \ldots, y^{(M)})^T$ by estimating and testing all relevant effects. For each predic-

tor x_i , we desire an estimation and testing procedure that will identify its associated subset of y. In particular, researchers often want to identify predictors that are important for multiple or all outcomes. We will call x_i a "multiple regulator" if it is associated with multiple outcomes, a terminology which we adapt from Peng et al. (2010). An example of what we call multiple regulation is known as pleiotropy in the genetics literature. Our goal of identifying multiple regulation is not to be confused with identifying predictors that are associated with any outcome. Association with any outcome is an active area of research, with two examples being global association tests and group-sparse regularization. Global tests provide a test for the relationship between x_i and the entire set y (Jiang and Zeng, 1995; He et al., 2013) and have been shown in some situations to have higher power than marginal tests to detect associations when x_i relates to multiple outcomes. Groupsparse methods, largely based on the group lasso (Yuan and Lin, 2006), use model selection to identify predictors that are relevant for any outcome (Turlach et al., 2005). These methods, while powerful and useful, do not address the question of which outcomes are relevant for each predictor and in general are unsuited for diverse outcomes that may contain censoring.

Here, we are particularly interested in identifying predictors that are relevant for multiple outcomes and inferring which subset of \mathbf{y} each of the x_j 's are associated with. There is a paucity of literature that addresses these specific questions. Under linear regression models, the remMap procedure (Peng et al., 2010) addresses such a question via variable selection by jointly penalizing both the L_1 and L_2 group

norms of a squared loss. Under generalized linear models, one could potentially modify the hierarchical lasso (Zhou and Zhu, 2010) procedure, originally proposed to handle grouped predictors with a single outcome, to address the multiple regulator problem. When making joint inference on a diverse set of outcomes, it is also desirable to put all effects on similar scales. A simple example of this idea can be found in (Schifano et al., 2013), where linear regression models were considered for multiple continuous outcomes and each outcome was scaled by its standard deviation. However, none of these methods are applicable to settings where y consists of a diverse set of outcomes whose scales may not be easily comparable to each other, especially when y may contain censored time-to-event variables. To accommodate modeling of multiple outcomes of different scales and/or type, we propose in this article the use of semiparametric transformation models which give all effects of \mathbf{x} on \mathbf{y} a similar interpretation. A liability thresholding version of such models can naturally model binary or ordinal outcomes.

Regardless of estimation technique, a multiple testing procedure is required to control error rates when identifying multiple regulation, which operates on the (potentially large) set of hypotheses $\{H_i^{(m)}: x_j \text{ unassociated with }$ $y^{(m)}_{j=1,\dots,p;m=1,\dots,M}$. Neither Peng et al. (2010) nor Zhou and Zhu (2010) tackles this issue. In general, multiple testing based on regularized estimation is challenging for two reasons. First, while many of the regularization procedures such as Zhou and Zhu (2010) established asymptotic oracle properties for their estimators—noninformative predictors can be detected with no uncertainty and their detection induces no additional variation in the estimation of the informative predictors (Fan and Li, 2001; Zou, 2006)—in finite samples those properties may be far from holding. Consequently, basing testing procedures on such asymptotic results may lead to inflated type I error in finite samples. Second, the estimators and hence their corresponding test statistics could be highly correlated from the regression fitting. Standard methods for controlling the familywise error rate (FWER), like the Bonferroni procedure, tend to be conservative in the presence of correlation, and they ignore the dependence structure in the data.

We propose a two-stage technique to both estimate the effects of \mathbf{x} on \mathbf{y} and identify multiple regulation while controlling error rates. In the first stage, we posit models to put all effects on the same scale, and we use regularization to induce sparsity in the estimated effects. To do this, we generalize the adaptive hierarchical lasso of Zhou and Zhu (2010) to handle the case of semiparametric models. In the second stage, we employ a stepdown procedure analogous to Romano and Wolf (2005) to identify multiple regulation while controlling error rates. Our two-stage method, entitled Sparse Multiple Regulation Testing (SMRT), is powerful for several reasons. First, our modeling strategy allows us to do estimation and make inference on outcomes that may be measured on completely different scales. Next, regularization enables us to more efficiently estimate both the null and non-null effects. The null effects are estimated as 0 with probability tending to 1 and the non-null effects are estimated with lower variability compared to unregularized estimators. Furthermore, the distributions of the estimates of null effects and the distributions of the estimates of non-null effects are distinctly separated through regularization, giving us more power to detect the non-null effects (see Figure 1 in Web Appendix A for an illustration from our simulations). Finally, our testing procedure can be specifically geared to detect associations with multiple outcomes.

However, it is generally challenging to perform testing based on regularized estimators since their distributions in finite samples cannot be approximated well by asymptotic results. We lay out permutation- and resampling-based procedures to better approximate the finite-sample distributions of the proposed test statistics and the regression parameter estimators. This enables us to properly control error rates for both hypothesis testing and interval estimation. Thus, in addition to providing the estimator β based on joint regularization, the main contributions of this article include providing resampling procedures to make joint inference about β and deriving the SMRT testing procedure to identify the subset of outcomes associated with each of the predictors. Our proposed estimation and testing procedures can account for the joint effects of the predictors and the correlation among both the predictors and the outcomes.

The rest of the article is organized as follows. In Section 2, we give an overview of SMRT. In Section 3, we discuss details regarding our sparse estimator, including its asymptotic properties and quantifying its variability. In Section 4, we discuss issues related to testing, including the asymptotic guarantee of familywise error control and practical approaches to finitesample error control. In Section 5, we apply our method to a genetic study of autoantibodies with the goal of identifying multiple regulators of autoimmunity. Simulation results which validate our method are provided in Section 6. And finally, in Section 7, we discuss implications and further directions.

2. Overview of SMRT

Suppose the data for analysis consists of *n* independent and identically distributed random vectors $\mathbb{V} = \{\mathbf{V}_i = (\mathbf{y}_i^T, \mathbf{x}_i^T)^T\}_{i=1,\dots,n}$ where $\mathbf{y}_i = (y_i^{(1)}, \dots, y_i^{(M)})^T$ are the *M* outcomes and $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^T$ are the *p* predictors for the *i*th subject. We first propose a unified modeling strategy for diverse \mathbf{y} by assuming that

$$P(y^{(m)} \le y \mid \mathbf{x}) = g^{(m)} \{ \mathbf{x}^{T} \boldsymbol{\beta}_{0}^{(m)} + h^{(m)}(y) \}, \qquad m = 1, \dots, M,$$
(1)

where $\boldsymbol{\beta}_{0}^{(m)}$ represents the unknown effect of \mathbf{x} on $y^{(m)}$, $h^{(m)}(\cdot)$ is an unspecified smooth, increasing function, and the link function, $g^{(m)}$, is given although the correlation structure of \mathbf{y} is left unspecified. For ease of presentation, we assume that \mathbf{y} is fully observed although the proposed method can easily accommodate censored outcomes. When $y^{(m)}$ is continuous, (1) is equivalent to

$$h^{(m)}(y^{(m)}) = -\mathbf{x}^{T} \boldsymbol{\beta}_{0}^{(m)} + \epsilon^{(m)},$$

with $P(\epsilon^{(m)} \le z \mid \mathbf{x}) = P(\epsilon^{(m)} \le z) = g^{(m)}(z).$ (2)

Generalized linear models for a binary or ordinal outcome can be written in the form of (1) and (2) by viewing the observed outcome as a thresholded version of a latent continuous outcome and $h^{(m)}$ as only defined at the threshold values, as previously suggested in the literature (Thomas et al., 1998, e.g.). Choice of $g^{(m)}$ determines the type of model being fit. For example, $g^{(m)}(x) = e^x/(1+e^x)$ corresponds to a proportional odds model for continuous $y^{(m)}$ and a logistic regression model if $y^{(m)}$ is binary. Models (1) and (2) have also been previously used to analyze censored survival outcomes (Cai et al., 2000; Zeng and Lin, 2007). The virtue of this approach is that the scale of the $\boldsymbol{\beta}^{(m)}$ will be comparable across m = 1, ..., M when the same or comparable $g^{(m)}(x)$ are used whether $y^{(m)}$ is continuous, discrete, or not fully observed because each marginal model has a similar form. For example if $g^{(m)}(x) = e^x/(1+e^x)$, then each $\beta_i^{(m)}$ has the interpretation of a log odds ratio regardless of whether $y^{(m)}$ is continuous, binary, ordinal, or censored.

To estimate $\boldsymbol{\beta}_{0}^{(m)}$, one may employ the nonparametric maximum-likelihood estimator (NPMLE) under model (1) (Murphy and Van der Vaart, 2000; Zeng and Lin, 2007) based on data observed on the *m*th outcome, $\mathbb{V}^{(m)} =$ $\{(\mathbf{y}_i^{(m)}, \mathbf{x}_i^T)^T\}_{i=1,\dots,n}$. Let $\mathcal{L}^{(m)}(\boldsymbol{\beta}^{(m)})$ denote the resulting profile log-likelihood (PLL) function corresponding to the NPMLE. It has been shown that under mild smoothness conditions, the profile likelihood can be treated as a regular likelihood, and the maximum PLL estimator $\widetilde{\boldsymbol{\beta}}^{(m)} = \operatorname{argmax}_{\boldsymbol{\beta}^{(m)}} \mathcal{L}^{(m)}(\boldsymbol{\beta}^{(m)})$ is regular and semiparametric efficient (Murphy and Van der Vaart, 2000). However, when p is not too small and $\beta_0 =$ $(\boldsymbol{\beta}_{0}^{(1)T},\ldots,\boldsymbol{\beta}_{0}^{(M)T})^{T}$ might be sparse, an improved estimator may be obtained by imposing regularization on the PLL. To do this, we simultaneously consider all M outcomes and obtain a sparse $\widehat{\boldsymbol{\beta}} = (\widehat{\boldsymbol{\beta}}^{(1)T}, \dots, \widehat{\boldsymbol{\beta}}^{(M)T})^T$ as the minimizer of the penalized sum of negative PLLs

$$-\sum_{m=1}^{M} \mathcal{L}^{(m)}(\boldsymbol{\beta}^{(m)}) + p_{\lambda,\mathbf{w}}(\boldsymbol{\beta}), \qquad (3)$$

with penalty function $p_{\lambda,\mathbf{w}}(\boldsymbol{\beta}) = \sum_{j=1}^{p} d_j + \lambda \sum_{m=1}^{M} \sum_{j=1}^{p} w_j^{(m)} |\alpha_j^{(m)}|$, with $\beta_j^{(m)} = d_j \alpha_j^{(m)}$, subject to $d_j \ge 0$. The penalty function $p_{\lambda,\mathbf{w}}(\cdot)$ was previously proposed in Zhou and Zhu (2010) for generalized linear models with grouped predictor variables. The tuning parameter λ controls the amount of regularization and weight $w_i^{(m)} = |\widetilde{\beta}_i^{(m)}|^{-1}$ is chosen to ensure oracle properties of $\hat{\boldsymbol{\beta}}$. Summing over the PLLs in (3) essentially imposes a working independence assumption across the outcomes (Liang and Zeger, 1986). Imposing the joint penalty $p_{\lambda,\mathbf{w}}(\boldsymbol{\beta})$ incorporates the potential for joint sparsity across all outcomes for some x_i 's. Setting $d_j = 0$ declares x_j to be noninformative for all outcomes or equivalently $\boldsymbol{\beta}_{0j} = (\beta_{0j}^{(1)}, \dots, \beta_{0j}^{(M)})^{T} = 0$; while setting $\alpha_{j}^{(m)} = 0$ suggests that $\beta_{0j}^{(m)} = 0$. We will show that $\hat{\boldsymbol{\beta}}$ possesses a sparsistency property, i.e., $P(\widehat{\beta}_{j}^{(m)} = 0 | \beta_{0j}^{(m)} = 0) \rightarrow 1$. This ensures desirable asymptotic properties for our testing procedures. We give further details regarding $\hat{\beta}$ and its asymptotic properties in Section 3. We now turn to the topic of testing.

2.1. Testing a Single Predictor x_i

In order to make inference on a single predictor, SMRT employs a stepdown procedure for x_j considering the M hypotheses $\mathcal{H}_j = \{H_j^{(m)} : \beta_{0j}^{(m)} = 0\}_{m=1,\dots,M}$ with alternative hypotheses denoted $\{\bar{H}_j^{(m)} : \beta_{0j}^{(m)} \neq 0\}_{m=1,\dots,M}$. To test $H_j^{(m)}$, we consider the statistic $t_j^{(m)} = n^{\frac{1}{2}} \left| \widehat{\beta}_j^{(m)} \right| / \widetilde{\sigma}_j^{(m)}$ and its reference distribution $\mathcal{T}_j^{(m)} = \{t_j^{*b(m)}\}_{b=1,\dots,B}$ which approximates the distribution of $t_j^{(m)} \mid H_j^{(m)}$ and can be obtained by, for example, resampling or permutation (see Section 4.2). We scale $\widehat{\beta}_j^{(m)}$ by $\widetilde{\sigma}_j^{(m)}$, which is an estimated standard error of $n^{\frac{1}{2}} (\widetilde{\beta}_j^{(m)} - \beta_{0j}^{(m)})$, since under $H_j^{(m)}, \widehat{\sigma}_j^{(m)} = \widehat{\operatorname{var}}\{n^{\frac{1}{2}} (\widehat{\beta}_j^{(m)} - \beta_{0j}^{(m)})\}^{1/2} \to 0$ and the null distribution of $n^{\frac{1}{2}} \widehat{\beta}_j^{(m)} / \widehat{\sigma}_j^{(m)}$ is difficult to approximate.

To test \mathcal{H}_j simultaneously, we order the test statistics $\mathbf{t}_j = (t_j^{(1)}, \ldots, t_j^{(M)})^T$ from largest to smallest, $t_j^{(r_1)} \geq t_j^{(r_2)} \geq \ldots \geq t_j^{(r_M)}$, and identify their corresponding hypotheses $H_j^{(r_1)}, \ldots, H_j^{(r_M)}$. Define for every $\Omega \subset \{1, \ldots, M\}$ the sup-statistic over Ω and its corresponding reference distribution: $s_j^{\Omega} = \max_{m \in \Omega} t_j^{(m)}$ and $\mathcal{S}_j^{\Omega} = \{\max_{m \in \Omega} t_j^{*b(m)}\}_{b=1,\ldots,B}$. Furthermore, denote the ψ th quantile of \mathcal{S}_j^{Ω} by $c_j^{\Omega}(\psi)$, which approximates the ψ th quantile of s_j^{Ω} under the null that $\{\beta_j^{(m)} = 0 : m \in \Omega\}$. We identify the subset of hypotheses to reject, denoted by \mathcal{R}_j , as follows.

- 1) Let $\Omega_1 = \{1, \ldots, M\}$. If $s_j^{\Omega_1} \leq c_j^{\Omega_1}(\psi)$, accept all hypotheses and stop. Otherwise, let $\mathcal{R}_j = \{r_1\}$ and continue. ...
- 1) Let $\Omega_l = \Omega_1 \setminus \mathscr{R}_j$. If $s_j^{\Omega_l} \leq c_j^{\Omega_l}(\psi)$, accept all hypotheses in $\{H_j^{(m)}\}_{m \in \Omega_l}$ and stop. Otherwise, let $\mathscr{R}_j = \mathscr{R}_j \cup \{r_l\}$ and continue. ...
- M) Let $\Omega_M = \{r_M\}$. If $s_j^{\Omega_M} \le c_j^{r_M}(\psi)$, accept $H_j^{(r_M)}$. Otherwise, let $\mathscr{R}_j = \mathscr{R}_j \cup \{r_M\}$.

The stepdown procedure for the simultaneous testing of \mathcal{H}_j then rejects all hypotheses in $\{H_j^{(m)}\}_{m \in \mathscr{R}_j}$ and concludes that x_j is associated with $\{y^{(m)}\}_{m \in \mathscr{R}_j}$. If the reference distribution and ψ are chosen such that the probability of making a type I error at each step is at most α :

$$P\left(s_{j}^{\Omega_{k}} > c_{j}^{\Omega_{k}}(\psi) \Big| \bigcap_{m \in \Omega_{k}} H_{j}^{(m)}\right) \leq \alpha, \tag{4}$$

for any k, then the FWER of the stepdown procedure—that is, the probability of making at least one false rejection over the set \mathcal{H}_j —is maintained at α . We discuss in detail issues relating to the choice of reference distribution and ψ in Section 4. We also describe how, regardless of the choices of reference distribution and ψ , the FWER is asymptotically 0 because $\hat{\beta}$ is sparsistent.

2.2. Multiple Regulation Testing

Now suppose scientific interest lies only with a predictor if it regulates at least k outcomes. That is, we only care to conclude that x_j is associated with $\{y^{(m)}\}_{m \in \mathcal{R}_j}$, for some $\mathscr{R}_j \subset \{1, \ldots, M\}$ if the number of rejections (i.e., the cardinality of \mathcal{R}_j) is at least k. Then, we can modify the testing procedure in the previous section to increase power to detect k-multiple regulators (kMRs) at the expense of being able to detect if x_j appears to be associated with fewer than k outcomes. The testing procedure proceeds by essentially skipping the first k - 1 steps in the previous section and only rejecting the first k - 1 hypotheses if any other hypotheses are rejected. Thus, we will either reject 0 hypotheses or k or more hypotheses. Throughout, when we refer to SMRT, we mean the combination of our sparse estimation technique and our multiple regulation testing procedure for a given k, with k = 1corresponding to the application of the test in the previous section.

We identify the subset of hypotheses to reject, denoted by \mathscr{R}_{j} , as follows: (i) let $\Omega_{1} = \{r_{k}, \ldots, r_{M}\}$. If $s_{j}^{\Omega_{1}} \leq c_{j}^{\Omega_{1}}(\psi)$, accept all hypotheses and stop. Otherwise, let $\mathscr{R}_{j} = \{r_{1}, \ldots, r_{k}\}$ and continue; (ii) let $\Omega_{2} = \{1, \ldots, M\} \setminus \mathscr{R}_{j}$. If $s_{j}^{\Omega_{2}} \leq c_{j}^{\Omega_{2}}(\psi)$, accept all hypotheses in $\{H_{j}^{(m)}\}_{m \in \Omega_{2}}$ and stop. Otherwise, let $\mathscr{R}_{j} = \mathscr{R}_{j} \cup \{r_{k+1}\}$ and continue. Steps 3 through M - k + 1 proceed as in the previous section.

As discussed in Section 4, the stepdown test with k > 1also has asymptotic FWER of 0. In addition to requiring (4), which we will call controlling the common type I error, we also require the control of a second type of error: *incorrectly* rejecting one of $\{H_j^{(m)}\}_{m=r_1,...,r_{k-1}}$ based on *correctly* rejecting $H_j^{(r_k)}$ in step one. We will call this a type I error by implication. Since the distribution of null effects gets shrunk dramatically toward 0 (see Figure 1 in Web Appendix A), it is unlikely for this type of error to occur in practice because it requires a test statistic corresponding to a null hypothesis to be larger than a test statistic from a rejected alternative hypothesis. We leave discussion of controlling the FWER for all predictors to Web Appendix A. The extension of the testing procedure for a single predictor is straightforward.

3. Inference about $\hat{\beta}$

We next detail the construction of $\hat{\boldsymbol{\beta}}$ as well as the asymptotic distribution for the zero and nonzero components, which is crucial for the validity of our estimator, confidence intervals, and proposed testing procedures. Estimation proceeds by minimizing (3). Now, since the profile log-likelihoods $\{\mathcal{L}^{(m)}\}_{m=1,\dots,M}$ are nonlinear functions without closed form in most cases, direct maximization of (3) may be numerically challenging, especially when p is not small. To reduce the computational complexity and enable the use of widely available software, we propose to take a quadratic expansion of $\mathcal{L}^{(m)}(\boldsymbol{\beta}^{(m)})$ in (3) similar to Zhang and Lu (2007) and Wang and Leng (2007). Specifically, we instead minimize

$$\|\widetilde{\mathbf{Y}} - \widetilde{\mathbb{X}}\boldsymbol{\beta}\|_{2}^{2} + p_{\lambda,\mathbf{w}}(\boldsymbol{\beta}), \qquad (5)$$

where $\widetilde{\mathbf{I}}^{(m)} = -\ddot{\mathcal{L}}^{(m)}(\widetilde{\boldsymbol{\beta}}^{(m)}), \ \ddot{\mathcal{L}}^{(m)}(\mathbf{b}) = \partial^2 \mathcal{L}^{(m)}(\mathbf{b})/\partial \mathbf{b} \partial \mathbf{b}^T, \ \widetilde{\mathbb{X}} = \text{diag}(\widetilde{\mathbf{A}}^{(1)}, \dots, \widetilde{\mathbf{A}}^{(M)}), \ \widetilde{\mathbf{Y}} = \widetilde{\mathbb{X}}\widetilde{\boldsymbol{\beta}} \text{ and } \widetilde{\mathbf{A}}^{(m)} \text{ is a symmetric half matrix of } \widetilde{\mathbf{I}}^{(m)} \text{ such that } \widetilde{\mathbf{I}}^{(m)} = \widetilde{\mathbf{A}}^{(m)}\widetilde{\mathbf{A}}^{(m)}.$ Computational simplifications and a full algorithm for fitting are discussed in Web Appendix D.

3.1. Asymptotic Theory

In this section, we present the properties of our proposed estimator $\hat{\boldsymbol{\beta}}$. It has the property of *sparsistency* in that it asymptotically sets truly null effects to exactly 0. Specifically, define \mathcal{A} and \mathcal{A}^c as indexing the nonzero and zero components of $\boldsymbol{\beta}_0$, respectively, where $\boldsymbol{\beta}_{\mathcal{A}}$ denotes the subvector of $\boldsymbol{\beta}$ corresponding to \mathcal{A} . Then, a *sparsistent* estimator $\hat{\boldsymbol{\beta}}$ is one that satisfies $P(\hat{\boldsymbol{\beta}}_{\mathcal{A}^c} = 0) \to 1$ as $n \to \infty$. Furthermore, our estimates of non-null effects are asymptotically normal and possess the *oracle property*, in that they are as efficient in the limit as if we knew which effects were truly null a priori. Let $\mathbf{I}_{\mathcal{A},\mathcal{B}}$ denotes the submatrix of \mathbf{I} corresponding to rows in \mathcal{A} and columns in \mathcal{B} .

In Web Appendix A, we show that for PLLs $\{\mathcal{L}^{(m)}(\boldsymbol{\beta}^{(m)})\}_{m=1,\dots,M}$ that satisfy certain regularity conditions (also listed in Web Appendix A), if $n^{-1}\sqrt{\lambda} = o_p(n^{-1/2})$, then there exists a root-*n* consistent local maximizer $\hat{\boldsymbol{\beta}}$ such that $P(\hat{\boldsymbol{\beta}}_{A^c} = 0) \rightarrow 1$ and $n^{\frac{1}{2}}(\hat{\boldsymbol{\beta}}_A - \boldsymbol{\beta}_{0,A}) \rightarrow N(0, \mathbf{I}_{A,A}^{-1}\boldsymbol{\Sigma}_{A,A}\mathbf{I}_{A,A}^{-1})$ in distribution, where $\boldsymbol{\Sigma}_{A,A} = \operatorname{cov}(\boldsymbol{\varphi}_{iA}(\boldsymbol{\beta}_0)), \, \boldsymbol{\varphi}_{iA}(\boldsymbol{\beta}_A)$ denotes the contribution of the *i*th subject to the profile score function for $\boldsymbol{\beta}_A$, $\mathbf{I} = \operatorname{diag}\{\mathbf{I}^{(1)}, \dots, \mathbf{I}^{(m)}\}$, and $\mathbf{I}^{(m)}$ is the limiting information matrix. This result, parallel to that given in Zhou and Zhu (2010), offers the promise of identifying null effects with probability approaching 1, while efficiently estimating non-null effects. From a testing perspective, it ensures that the type I error of SMRT for any k decreases to 0 as $n \to \infty$.

3.2. Estimating the Variability in $\widehat{\beta}$

The asymptotic results on $\hat{\boldsymbol{\beta}}$ suggest that we are as efficient in the limit as if we knew which parameters were truly 0 from the outset. However, in finite samples the added variability due to estimating \mathcal{A}^c may not be negligible, and hence relying on the asymptotic result will underestimate the variability in $\hat{\boldsymbol{\beta}}$. To better approximate the finite-sample distribution, we propose a perturbation resampling procedure to estimate the distribution of $n^{\frac{1}{2}}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)$. This procedure, by accounting for the variability in estimating \mathcal{A}^c , provides a more precise estimate of the variability in $\hat{\boldsymbol{\beta}}$ and maintains the correlation structure in $\hat{\boldsymbol{\beta}}$.

We generate a resampled counterpart of $\widehat{\boldsymbol{\beta}}$, denoted by $\widehat{\boldsymbol{\beta}}^*$, in two steps. We first generate $\widetilde{\boldsymbol{\beta}}^*$, a resampled version of $\tilde{\boldsymbol{\beta}}$, by either perturbing the profile likelihood or directly perturbing the influence function corresponding to $\hat{\boldsymbol{\beta}}$. In essence, each perturbation is achieved by multiplying G_i to the likelihood contribution from the *i*th subject, where the positive perturbation variables $\{G_i\}$ are generated independently with mean 1 and variance 1. Then, we minimize our objective function (5) using $\widetilde{\boldsymbol{\beta}}^*$ in place of $\widetilde{\boldsymbol{\beta}}$, yielding resampled estimates $\widehat{\boldsymbol{\beta}}^*$. Similar resampling procedures have been proposed for making inference with a wide range of standard objective functions without regularization (Tian et al., 2007; Uno et al., 2007, e.g.) and recently extended to accommodate L_1 type regularized estimators (Minnier et al., 2011). Here, we propose such a resampling procedure to both account for the potential correlation among the outcomes and better approximate the finite-sample behavior of hierarchically regularized estimators.

In Web Appendix B, we detail the perturbation procedure and establish its asymptotic properties, which are parallel to those for $\hat{\boldsymbol{\beta}}$. A key feature of the resampled $\hat{\boldsymbol{\beta}}^*$ is that $n^{\frac{1}{2}} \left(\hat{\boldsymbol{\beta}}_{\mathcal{A}}^* - \hat{\boldsymbol{\beta}}_{\mathcal{A}} \right) | \mathbb{V}$ has the same limiting distribution as $n^{\frac{1}{2}} \left(\hat{\boldsymbol{\beta}}_{\mathcal{A}}^* - \hat{\boldsymbol{\beta}}_{\mathcal{A}} \right) | \mathbb{V}$ has the same limiting distribution of $\hat{\boldsymbol{\beta}}$ for a given dataset, we may generate a large number of $\hat{\boldsymbol{\beta}}^*$ s, $\left\{ \hat{\boldsymbol{\beta}}^{*b} \right\}_{b=1,\dots,B}$ for some suitably large *B*. To construct a confidence interval (CI) for a specific $\beta_j^{(m)}$, one may estimate the standard error of $\hat{\beta}_j^{(m)}$ as $\hat{\sigma}_j^{(m)}$ the empirical standard error of its perturbed realizations, $\left\{ \hat{\beta}_j^{*b(m)} \right\}_{b=1,\dots,B}$. A $100(1-\alpha)\%$ level confidence interval can then be constructed based on the normal confidence interval $\hat{\beta}_j^{(m)} \pm \mathcal{Z}_{1-\alpha/2} \hat{\sigma}_j^{(m)}$ or alternatively the lower and upper $\alpha/2$ percentiles of $\left\{ \hat{\beta}_j^{*b(m)} \right\}_{b=1,\dots,B}$.

3.3. Tuning

SMRT involves a large number of minimizations and tuning parameter selections. It is thus not feasible to select λ using time-consuming methods such as cross-validation. We propose a modified BIC criteria: $\lambda = \operatorname{argmin}_{\lambda} \left(\| \widetilde{\mathbf{Y}} - \widetilde{\mathbb{X}} \boldsymbol{\beta}_{\lambda} \|_{2}^{2} + \min\{n^{0.1}, \log n\} \operatorname{df}_{\lambda} \right)$, where $\boldsymbol{\beta}_{\lambda}$ is the minimizer of (5) corresponding to λ and $\operatorname{df}_{\lambda}$ is the number of nonzero entries in $\boldsymbol{\beta}_{\lambda}$. In small and moderate sample sizes, $n^{0.1}$ is much smaller than $\log n$ and is used here. However, when *n* becomes large $\log n$ may be preferred. Wang and Leng (2007) showed that this BIC criteria (with either $\log n$ or $n^{0.1}$) satisfies the rate requirements for a standard adaptive LASSO-type penalty. Similar arguments can be used to justify the rate for the adaptive hierarchical LASSO-type penalty used in (5).

4. Testing

In this section, we show that the FWER of our testing procedure is asymptotically 0 because of the sparsistency of $\hat{\beta}$. We also discuss in more detail the choice of reference distribution.

4.1. Properties of SMRT

One of the main results of this article is that, given a suitably estimated β , the FWER of our stepdown procedure approaches 0 as $n \to \infty$ for any k regardless of the reference distribution or what quantile ψ we use to determine the cutoff for rejection. Specifically, we show in Web Appendix C that if $\boldsymbol{\beta}$ is sparsistent, then for every j and Ω , $P\left(s_{j}^{\Omega}>c_{j}^{\Omega}(\psi)\middle|\bigcap_{m\in\Omega}H_{j}^{(m)}\right)\to 0$ as $n\to\infty$, and SMRT has an asymptotic FWER of 0, for any reference distribution, k, and ψ . The result follows from showing that common type I errors and type I errors by implication both occur with probability tending to 0. With regard to common type I errors, under a given null $H_j^{(m)}$, the test statistic $t_j^{(m)}$ is estimated at exactly 0 with probability tending to 1 and, under the composite null $\bigcap_{m\in\Omega} H_j^{(m)}$, s_j^{Ω} tends to 0 as well. Thus, we cannot reject $\bigcap_{m \in \Omega} H_j^{(m)}$, regardless of the value of $c_j^{\Omega}(\psi)$, and therefore common type I errors will occur with probability approaching 0 as $n \to \infty$. The other potential source of type I error occurs for k > 1 when incorrectly rejecting $H_i^{(r_{k'})}$ based

on correctly rejecting $H_j^{(r_k)}$, k' < k. However, this sort of type I error will only occur if the test statistic for a null hypothesis (which is tending to 0) is larger in magnitude than the test statistic for an alternative hypothesis, which is of course impossible asymptotically.

While the foregoing result shows that the asymptotic behavior of SMRT is ensured by the sparsistency of $\hat{\beta}$, of course in finite samples choice of reference distribution and ψ is paramount in maintaining the desired error rate. Maintaining the FWER at approximately α can be ensured by choosing the reference distribution and ψ such that the probability of making a type I error at each step of the testing procedure is maintained at approximately α .

4.2. Choosing a Reference Distribution

As discussed in the previous section, any reference distribution $\mathcal{T}_{j}^{(m)} = \{t_{j}^{*_{b}(m)}\}_{b=1,\dots,B}$ will provide asymptotic control of the FWER by virtue of sparsistency of the estimator $\hat{\boldsymbol{\beta}}$. We explore resampling- and permutation-based reference distributions. The resampling-based reference distribution is based on $t_j^{*_b(m)} = n^{\frac{1}{2}} \left| \widehat{\beta}_j^{*_b(m)} - \widehat{\beta}_j^{(m)} \right| / \widetilde{\sigma}_j^{(m)}$. Simulation results suggest that, although resampling provides good approximation to the finite-sample distribution of $\hat{\boldsymbol{\beta}}_{\mathcal{A}}$, it tends to over-estimate the variability of $\hat{\boldsymbol{\beta}}_{A^c}$ (see Figure 1). As an alternative, we consider a permutation-based reference distribution with $t_i^{*_b(m)}$ based on an estimate of $\boldsymbol{\beta}_0$ from a dataset where $y^{(m)}$ is permuted. See Web Appendix E for further details about the procedure and Section 6 for simulation results. Numerical results suggest that, the permutation-based reference distribution does a better job of approximating the finite-sample null distribution of $t_i^{(m)}$, as shown in Figure 1.

4.3. Choosing ψ

To control the FWER at α -level, it seems reasonable to choose $\psi = 1 - \alpha$. This ensures that, for a suitable reference distribution and *n* large enough, $P\left(s_{j}^{\Omega} > c_{j}^{\Omega}(\psi) \middle| \bigcap_{m \in \Omega} H_{j}^{(m)}\right) \lessapprox \alpha$, for any Ω , which, along with negligible type I error by implication, will give approximate FWER control. In light of the fact that all type I errors are tending to probability 0, one could obtain improved power by choosing $\psi < 1 - \alpha$, while maintaining the level at α . This is particularly important if using the resampling-based reference distribution. One could use another layer of permutation or resampling to estimate the smallest ψ that would still maintain the level α . However, that requires computing a large number of permutations/resamples for each of the B members of the reference distribution, which becomes prohibitively computationally demanding quickly. Computing a suitable $\psi < 1 - \alpha$ is a topic of future research.

5. Genetic Study to Identify Shared Autoimmune Risk Loci

We apply SMRT to a study of shared autoimmunity with the goal of identifying genetic markers associated with four autoantibodies: anti-nuclear antibodies (ANA), anti-cyclic citrullinated protein (CCP) antibodies, anti-transglutaminase (TTG) antibodies, and anti-thyroid peroxidase antibodies (TPO). These four autoantibodies are, respectively, markers



Figure 1. Simulation-based empirical and estimated distribution of null effects. Empirical null distribution of $\hat{\beta}_{j}^{(m)}$ (labeled "Empirical") agrees closely in the tails with the permutation-based estimate (labeled "Permutation estimated"), while the resampling-based estimate (labeled "Resampling estimated") overestimates the density in the tails.

for four autoimmune diseases (ADs): systemic lupus erythematosus (SLE), rheumatoid arthritis (RA), celiac disease, and autoimmune thyroid disease. The genetic markers consists of 67 single-nucleotide polymorphisms (SNPs) previously published as potential risk markers for these four ADs. Discovering which SNPs regulate multiple ADs can aid in understanding potential shared pathways or etiology of these diseases (Zhernakova et al., 2009). While it is rare for an individual to have multiple ADs, multiple autoantibodies can be present in individuals predisposed to having the multiple ADs even in the absence of the disease phenotypes. Here, we consider the autoantibodies markers for subjects at higher risk for the ADs.

The study cohort includes 1265 individuals of European ancestry with RA identified through electronic medical records at Partners Healthcare (Liao et al., 2010). Due to a limit of quantification, the antibody measurements are highly unreliable when the values are either very low or very high. A convenient approach to incorporating such limitations is by assuming a marginal proportional odds model and truncating the observations at the limit of quantification. Hence, $\beta_{0j}^{(m)}$ still has the interpretation of being a log odds ratio (OR).

Results for the autoantibody data are summarized in Figure 2. Figure 2a shows results for the sparse estimation step. In the figure, SNPs are denoted along the *y*-axis, and outcomes are denoted along the *x*-axis. Color of the tile indicates the OR estimate, with darker colors indicating stronger association. In order to measure the strength of association with respect to the FWER, we provide adjusted *p*-values as the smallest α such that the test would reject while controlling the FWER for the SNP at α . Figure 2b shows this *p*-value for each test.

Due to the large number of hypotheses, we do not have sufficient power to detect multiple regulation while simultaneously controlling the FWER across all SNPs. Taking a less conservative view, if we control the FWER at the SNP level, five SNPs show some evidence of multiple regulation at $\alpha = 0.1$. The two strongest associations were with rs2187668 and rs3129860. Having previously shown associations to SLE (Taylor et al., 2011) and celiac (van Heel et al., 2007), rs2187668 was estimated to be related to the autoantibodies for those diseases at OR = 1.45 (*p*-value 0.005) for ANA and OR = 1.62 (*p*value 0.005) for TTG, as well as to CCP (OR = 0.78, *p*-value 0.05). This SNP is in the MHC region, which is known to affect immune function. Similarly, rs3129860, also in the MHC region, which had previously shown an association to SLE (Taylor et al., 2011), here demonstrated an association to ANA (OR = 1.28, *p*-value 0.05), CCP (OR = 1.50, *p*-value 0.003), and TPO (OR = 1.30, *p*-value 0.05).

6. Simulation Results

We ran simulations to assess the performance of our point and interval estimation procedures as well as SMRT. We loosely based our simulations on the autoantibody dataset, allowing the relationship between **x** and **y** to be specified by a proportional odds model. We considered sample sizes of 150, 250, and 500 and ran 1000 simulations for each sample size. For each simulation, 1000 resampled $\hat{\boldsymbol{\beta}}^*$ s were generated.

We set the number of predictors of interest p to be 30 and the number of outcomes M to be 4. Covariates \mathbf{x} took values in $\{0, 1, 2\}$ with probability $\{\pi^2, 2\pi(1 - \pi), (1 - \pi)^2\}$ where $\pi = 0.15$. Outcomes \mathbf{y} were generated according to the marginal proportional odds model, conditional on \mathbf{x} . We allowed correlation in \mathbf{y} , which was accomplished by first generating correlated normal random variables $\mathbf{z}_i \sim N_4(\mathbf{0}, \Sigma)$ where $\Sigma = 0.85I + 0.15\mathbf{11}^T$ is exchangeable. Then, let $\mathbf{u}_i = \Phi(\mathbf{z}_i)$ for Gaussian distribution function $\Phi(\cdot)$, and finally $\mathbf{y}_i = \exp(\mathbf{x}_i \boldsymbol{\beta}_0 + \boldsymbol{\epsilon}_i)$ where $\boldsymbol{\epsilon}_i = \log(\frac{u_i}{1-u_i}) \sim$ logistic. For computational simplicity, we discretized \mathbf{y} into 10 levels of roughly equal sizes according to deciles. The only change when discretizing is to the number of locations at which $h^{(m)}$ is



Figure 2. Results for autoantibody data. SNPs are listed on the *y*-axis, and autoantibodies are listed on the *x*-axis. (a) Sparse odds ratio effect estimates. Darker colors indicate larger magnitudes, and white indicates no estimated association. (b) Adjusted *p*-values. Darker color indicates smaller *p*-value and more evidence against the null hypothesis of no association.

estimated. In practice, this is not an issue (note that we did not discretize in the data analysis), but for the purposes of simulation it was a moderate speed-up with little information loss.

The relationship between \mathbf{x} and \mathbf{y} is defined by

$$(\boldsymbol{\beta}_{0}^{(1)},\ldots,\boldsymbol{\beta}_{0}^{(M)})_{30\times4} = \begin{pmatrix} \mathbf{1}_{20} & \frac{1}{2}_{16} & \mathbf{1}_{12} & \frac{1}{2}_{8} \\ \mathbf{0}_{10} & \mathbf{0}_{14} & \mathbf{0}_{18} & \mathbf{0}_{22} \end{pmatrix}_{30\times4}$$

where $\mathbf{1}_k$ is a $k \times 1$ vector of ones, $\mathbf{0}_k = 0 \times \mathbf{1}_k$, and $\frac{1}{\mathbf{2}_k} = \frac{1}{2} \times \mathbf{1}_k$. This configuration indicates that there are eight predictors related to all four outcomes, four related to just the first three outcomes, four related to just the first two outcomes, and four related to just the first outcome. The remaining 10 predictors are null, unrelated to any outcome. We also see that associations to outcomes $y^{(2)}$ and $y^{(4)}$ are

weak, so we would expect there to be less power to detect those effects.

6.1. Estimation

We first demonstrate that our point and interval estimation procedures perform well in finite samples. Figure 3 (top panel) shows the average bias in $\hat{\boldsymbol{\beta}}$ and $\tilde{\boldsymbol{\beta}}$ across simulations, plotted according to true effect size $\boldsymbol{\beta}_0$ and sample size. The regularized $\hat{\boldsymbol{\beta}}$ exhibits much smaller bias than the unregularized $\tilde{\boldsymbol{\beta}}$ for all sample sizes and effect sizes. Particularly at smaller sample sizes, regularization substantially reduces the bias in the estimator.

In Figure 3 (middle panel), we plot the average bias in SE estimates obtained based on our proposed resampling procedures as well as those based on the asymptotic variance. Both the asymptotic SE estimate and the resampling-based



Figure 3. Average performance of point estimates, standard errors, and confidence intervals across 1000 simulations at sample sizes of n = 150, 250, 500. All quantities are aggregated over $\beta_{0j}^{(m)}$ and plotted against $\beta_{0j}^{(m)}$. Top panel: average estimated bias in regularized $\hat{\beta}_{j}^{(m)}$ and unregularized $\tilde{\beta}_{j}^{(m)}$. Middle panel: average estimated bias of estimates of $\sigma_{j}^{(m)}$, the standard error of $\hat{\beta}_{j}^{(m)}$, comparing resampling-based estimates to asymptotic estimates. Bottom panel: 95% CI coverage comparing asymptotic, resampling-based normal, and resampling-based quantile CIs.

one $\widehat{\sigma}_{j}^{(m)}$ overestimate the variability in $\widehat{\beta}_{j}^{(m)}$ when $\beta_{0j}^{(m)} = 0$, but $\widehat{\sigma}_{j}^{(m)}$ more closely approximates $\sigma_{j}^{(m)}$. When $\beta_{0j}^{(m)} \neq 0$, the asymptotic SE tends to underestimate the true variability, while $\widehat{\sigma}_{j}^{(m)}$ approximates it well.

coverage levels of asymptotic-based CIs are as low as 78% for nonzero effects and remain lower than the nominal level even when n = 500. Hence in practice, we recommend the quantile-based CIs.

We examine CI coverage in the bottom panel of Figure 3 and see that underestimating the SEs leads to poor 95% CI coverage levels for the normal-based CI methods, based on $\tilde{\sigma}_{j}^{(m)}$ and $\hat{\sigma}_{j}^{(m)}$. Resampling-based quantile 95% CIs have good coverage for all values of $\beta_{0j}^{(m)}$ and all sample sizes. The

6.2. Testing

In the following sections, we examine the performance of SMRT. We first characterize the performance of our procedure with k = 1 when testing is performed for each predictor individually considering both the resampling-based and the



Figure 4. Threshold ψ (y-axis) plotted against its associated empirical rejection rate (x-axis) for the marginal test of $H_j^{(m)}$ across 1000 simulations, with color denoting the magnitude of $\beta_{j0}^{(m)}$ and linetype indicating whether regularization was used. The value on the y-axis ψ_r represents the threshold at which the empirical type I error was controlled for the regularized test, and ψ_u represents the threshold at which the empirical type I error was controlled for the unregularized test. Results for effects of the same magnitude are averaged for ease of presentation.

permutation-based reference distribution. Then, we consider testing with k > 1 and controlling error rates for all predictors.

6.3. Resampling-Based Reference Distribution

We briefly demonstrate the gains in power possible by using the resampling-based reference distribution. For ease of presentation, we demonstrate the performance of the testing procedure for the marginal test of $H_j^{(m)}$ with and without regularization. Results for the full stepdown procedure are similar.

Figure 4 demonstrates the power gain possible when using the regularized estimator with the resampling-based reference distribution. The plot shows the threshold necessary to obtain a given rejection rate. The ideal threshold maintains the rejection rate for null effects $(\beta_{0j}^{(m)} = 0)$ at a given level, say $\alpha = 0.05$, indicated by the vertical-dashed line. That threshold that maintains the type I error for the regularized estimator, indicated by ψ_r in the plot, is much lower than the threshold for the unregularized estimator, indicated by $\psi_u = 1 - \alpha = 0.95$ in the plot. Furthermore, the power to detect weak effects $(\beta_{0j}^{(m)} = 0.5)$ using the regularized estimator at ψ_r is 58% compared to 52% using the unregularized estimator at ψ_u , while the power to detect strong effects is similar. Thus, if one could select ψ_r adaptively, it appears that large power gains could be observed by using regularization. Due to its computational burden, however, we did not pursue this method further in our simulations.

6.4. Permutation-Based Reference Distribution

We pursue a more rigorous study of SMRT using the permutation-based reference distribution mentioned in Section 4.2 and described in detail in Web Appendix E. To demonstrate the role of regularization in improving testing, we compare SMRT to an identical testing procedure based on the unregularized $\tilde{\boldsymbol{\beta}}$, named MRT. We use the permutation-based reference distribution for both SMRT and MRT and take $\psi = 1 - \alpha$. To demonstrate the advantages of the step-down method, we compare to a single-step procedure, denoted as Sup, where we reject all $H_j^{(m)}$ for which $t_j^{(m)} > c_j^{\Omega_1}(\psi)$ where $\Omega_1 = \{1, \ldots, M\}$. Finally, we compare to the Bonferroni adjustment.

When controlling the FWER at $\alpha = 0.05$ for each x_j using the basic test, SMRT and MRT performed similarly in controlling FWER. The average empirical FWER was .046, .052, and .055 at n = 150, 250, 500, respectively for SMRT. The corresponding average FWER for MRT was 0.042, 0.049, 0.054, at those respective sample sizes. The more conservative Sup test had average FWERs of .041, .044, and .043, respectively, and the even more conservative Bonferroni .028, .026, .021.

In terms of power, SMRT dominates all other test procedures. Figure 5 depicts the power to detect non-null effects at n = 250 (other sample sizes show similar relative performances, with SMRT performing relatively better as sample size decreases). Possible rejections are listed across the bottom, and results are arranged according to how many outcomes the predictor is actually associated with. The figure shows that SMRT is uniformly more powerful than MRT, Bonferroni, and Sup, with the differences becoming more apparent in identifying multiple regulation.

Results for controlling the FWER by applying SMRT to all predictors and all hypotheses were qualitatively similar. All methods maintained the nominal level of the test, and SMRT obtained higher power than MRT, Sup, and Bonferroni at all sample sizes. When we apply SMRT to each x_j with k = 2, the average FWER for SMRT decreases to .020, .027, and .033 at n = 150, 250, 500. Taking k = 3 or k = 4 sees a further reduction to FWERs.

6.5. Superiority of Joint Analysis over Marginal Models

In this section, we demonstrate the advantages of performing a joint analysis for the detection of multiple regulation. We compare our estimator $\hat{\beta}$ to the estimator obtained by fitting each marginal model individually using L_1 penalization, which we will denote β^{\dagger} . The joint analysis improves our ability to detect multiple regulation, with the improvement over β^{\dagger} increasing with the number of outcomes a predictor is associated with. For example, when n = 500, for the eight predictors associated with all four outcomes, the power to detect association with $y^{(4)}$ and $y^{(2)}$ increased from 57 to 67% and from 60 to 66%, respectively, when using SMRT based on the joint penalty as opposed to marginal models, with a negligible increase for $y^{(3)}$ and $y^{(1)}$. For predictors of three outcomes,



Figure 5. Power to detect non-null effects across 1000 simulations at sample size n = 250 and level $\alpha = 0.05$. Each plot indicates how many outcomes the predictors tested are associated with. For example, the top left plot corresponds to predictors with strong association to $y^{(1)}$, $\beta_{0j}^{(1)} = 1$ and weak association to $y^{(2)}$, $\beta_{0j}^{(2)} = 0.5$. Tests are listed on the *x*-axis. Power is indicated on the *y*-axis. Power estimates are aggregated over all estimates that share the same effect sizes. To take a couple of examples, the bar corresponding to "y1" in the figure corresponds to power to reject $H_j^{(1)}$, and the bar corresponding to "y1/y2/y3" in the figure corresponds to power to reject each of $H_j^{(1)}$, $H_j^{(2)}$, $H_j^{(3)}$ simultaneously.

SMRT based on $\hat{\boldsymbol{\beta}}$ increased the power for association with $y^{(2)}$ to 61% from 57% for $\boldsymbol{\beta}^{\dagger}$. For predictors of two outcomes, $\hat{\boldsymbol{\beta}}$ had 58% power in detecting effects associated with $y^{(2)}$ compared to 54% for $\boldsymbol{\beta}^{\dagger}$. Furthermore, $\hat{\boldsymbol{\beta}}$ is much better at eliminating completely noninformative predictors by estimating all of their effects at exactly 0. Using joint estimation, $\hat{\boldsymbol{\beta}}$ eliminates null predictors completely 52% of the time, while the rate is only 23% using marginal models, when n = 500. The relative performance patterns are similar for n = 150 and 250.

7. Discussion

We have proposed a framework for testing and estimation across a diverse set of outcomes, with the explicit goal of identifying predictors for multiple outcomes. This framework allows the combination of information across continuous, semi-continuous, and discrete outcomes while maintaining control of the FWER. We have extended existing sparse regression methods for identifying multiple regulation to the complex scenario when the components of \mathbf{y} may be on very different scales or not completely observed. We have proven the asymptotic properties of this estimator and shown that one can use resampling to estimate its variability. We have, finally, provided a testing framework for identifying multiple regulation and demonstrated that the properties of the estimator ensure that the testing procedure has asymptotic FWER of 0.

While we rely on the sparsistency properties of our estimator, other penalty functions could potentially accomplish similar results to the hierarchical penalty we proposed. As long as sparsistency holds and a suitable finite-sample reference distribution can be obtained, e.g., through permutation or resampling, other penalty functions could be worth exploring. For simplicity, we used a working independence assumption to combine the profile log-likelihoods of multiple outcomes. But when the outcomes are not independent, incorporating information about the covariance in \mathbf{y} can improve efficiency (Liang and Zeger, 1986). A further advantage to using the quadratic approximation to $\mathcal{L}^{(m)}$ in (5) instead of the profile log-likelihood itself (besides computational tractability) is that we can incorporate covariance information about **y** through the initial estimate $\boldsymbol{\beta}$. If the (unpenalized) initial estimate $\boldsymbol{\beta}$ is estimated in a way that gains efficiency by taking correlation in \mathbf{v} into account, then that increase in efficiency will be propagated into our estimation of $\boldsymbol{\beta}$.

Due to the fine-grained nature of multiple regulation analysis and the complexity of dealing with diverse \mathbf{y} , SMRT may not be preferred in genome-wide or other very highdimensional data where discovery is of primary importance. Rather than using SMRT to discover novel markers, we suggest using it to validate known markers. Global tests for all outcomes (Jiang and Zeng, 1995; He et al., 2013) provide better power to discover unknown risk markers. Theoretically, the convergence of our estimators cannot be guaranteed jointly unless the number of predictors p and outcomes M is finite. Thus, we require M not to be too large compared to the sample size. Practically speaking, the computational complexity of the estimation procedure grows with M. A brief simulation yielded average run times of 0.4,1.7, 4.3, 11.0, 21.2, 44.9, and 343.8 seconds for M = 4, 8, 12, 16, 20, 25, and 50, respectively, at n = 500, with results being quite similar with n = 150.

Finally, we have focused on FWER as the error rate of primary interest throughout this article, but it could be of interest in some testing situations to employ less restrictive error control, especially when the number of tests grows large and signals are weak. We could easily extend SMRT with k = 1 to include more generalized error rates, such as k-FWER or the false discovery proportion, as in (Romano et al., 2010), and testing when k > 1 could be adapted in that direction as well.

8. Supplementary Materials

Web Appendices A, B, C, D, and E referenced in Sections 1, 2, 3, and 4 are available with this article at the *Biometrics* website on Wiley Online Library. An R package implementing the proposed estimation, perturbation, and testing procedures is also available at the *Biometrics* website and may also be accessed at github.com/denisagniel/smrtr.

References

- Cai, T., Wei, L., and Wilcox, M. (2000). Semiparametric regression analysis for clustered failure time data. *Biometrika* 87, 867– 878.
- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* 96, 1348–1360.
- He, Q., Avery, C. L., and Lin, D.-Y. (2013). A general framework for association tests with multivariate traits in large-scale genomics studies. *Genetic Epidemiology* **37**, 759–767.
- Jiang, C. and Zeng, Z.-B. (1995). Multiple trait analysis of genetic mapping for quantitative trait loci. *Genetics* 140, 1111– 1127.
- Liang, K.-Y. and Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika* 73, 13–22.
- Liao, K. P., Cai, T., Gainer, V., Goryachev, S., Zeng-treitler, Q., Raychaudhuri, S., et al. (2010). Electronic medical records for discovery research in rheumatoid arthritis. *Arthritis Care* & Research 62, 1120–1127.
- Minnier, J., Tian, L., and Cai, T. (2011). A perturbation method for inference on regularized regression estimates. *Journal of the American Statistical Association* **106**, 1371–1382.
- Murphy, S. A. and Van der Vaart, A. W. (2000). On profile likelihood. Journal of the American Statistical Association 95, 449–465.
- Peng, J., Zhu, J., Bergamaschi, A., Han, W., Noh, D. Y., Pollack, J. R., and Wang, P. (2010). Regularized multivariate regression for identifying master predictors with application to integrative genomics study of breast cancer. *The Annals of Applied Statistics* 4, 53–77.
- Romano, J. P. and Wolf, M. (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association* **100**, 94–108.
- Romano, J. P., and Wolf, M. (2010). Balanced control of generalized error rates. *The Annals of Statistics* 38, 598–633.
- Schifano, E. D., Li, L., Christiani, D. C., and Lin, X. (2013). Genome-wide association analysis for multiple continuous secondary phenotypes. *The American Journal of Human Genetics* 92, 744–759.
- Solovieff, N., Cotsapas, C., Lee, P. H., Purcell, S. M., and Smoller, J. W. (2013). Pleiotropy in complex traits: Challenges and strategies. *Nature Reviews Genetics*.
- Taylor, K. E., Chung, S. A., Graham, R. R., Ortmann, W. A., Lee, A. T., Langefeld, C. D., et al. (2011). Risk alleles for systemic lupus erythematosus in a large case-control collection and associations with clinical subphenotypes. *PLoS Genetics* 7, e1001311.
- Thomas, R., Have, T., Kunselman, A. R., Pulkstenis, E. P., and Landis, J. R. (1998). Mixed effects logistic regression models for longitudinal binary response data with informative dropout. *Biometrics* 54, 367–383.
- Tian, L., Cai, T., Goetghebeur, E., and Wei, L. (2007). Model evaluation based on the sampling distribution of estimated absolute prediction error. *Biometrika* 94, 297–311.
- Turlach, B. A., Venables, W. N., and Wright, S. J. (2005). Simultaneous variable selection. *Technometrics* 47, 349–363.
- Uno, H., Cai, T., Tian, L., and Wei, L. (2007). Evaluating prediction rules for t-year survivors with censored regression models. *Journal of the American Statistical Association* 102, 527–537.
- van Heel, D. A., Franke, L., Hunt, K. A., Gwilliam, R., Zhernakova, A., and Inouye, M., (2007). A genome-wide association study for celiac disease identifies risk variants in the region harboring IL2 and IL21. *Nature Genetics* **39**, 827–829.

- Wang, H. and Leng, C. (2007). Unified lasso estimation by least squares approximation. Journal of the American Statistical Association 102, 1039–1048.
- Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B* 68, 49–67.
- Zeng, D. and Lin, D. (2007). Maximum likelihood estimation in semiparametric regression models with censored data. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 69, 507–564.
- Zhang, H. H. and Lu, W. (2007). Adaptive lasso for cox's proportional hazards model. *Biometrika* 94, 691–703.
- Zhernakova, A., van Diemen, C. C., and Wijmenga, C. (2009). Detecting shared pathogenesis from the shared genetics of immune-related diseases. *Nature Reviews Genetics* 10, 43– 55.
- Zhou, N. and Zhu, J. (2010). Group variable selection via a hierarchical lasso and its oracle property. arXiv preprint arXiv:1006.2871.
- Zou, H. (2006). The adaptive lasso and its oracle properties. Journal of the American statistical association 101, 1418–1429.

Received January 2015. Revised November 2015. Accepted December 2015.