

Simultaneous detection and estimation of trait associations with genomic phenotypes

Jean Morrison* †, Noah Simon†, Daniela Witten‡

† *Department of Biostatistics*

‡ *Departments of Statistics and Biostatistics*

University of Washington, Seattle WA, 98195, USA

jeanm5@uw.edu

ABSTRACT

Genomic phenotypes, such as DNA methylation and chromatin accessibility, can be used to characterize the transcriptional and regulatory activity of DNA within a cell. Recent technological advances have made it possible to measure such phenotypes very densely. This density often results in spatial structure, in the sense that measurements at nearby sites are very similar.

In this paper, we consider the task of comparing genomic phenotypes across experimental conditions, cell types, or disease subgroups. We propose a new method, Joint Adaptive Differential Estimation (JADE), which leverages the spatial structure inherent to genomic phenotypes. JADE simultaneously estimates smooth underlying group average genomic phenotype profiles, and detects regions in which the average profile differs between groups. We evaluate JADE's performance in several biologically plausible simulation settings. We also consider an application to the detection of regions with differential methylation between mature skeletal muscle cells, myotubes and myoblasts.

*To whom correspondence should be addressed.

1. INTRODUCTION

During the past decade, it has become possible to measure *genomic phenotypes*, or local properties of DNA such as DNA methylation, chromatin accessibility, copy number variation, and histone modification. These genomic phenotypes can be measured very densely, and in some cases even at single-nucleotide resolution. For example, DNA methylation proportion can be measured at nearly every CpG using bisulfite sequencing, and chromatin accessibility can be measured at every nucleotide with DNase-seq. Neighboring nucleotides may belong to the same functional unit. Thus, genomic phenotypes often have similar values at nearby genomic positions.

Unlike the DNA sequence itself, genomic phenotypes may change dynamically as a function of cellular activity, developmental stage, and environment. This motivates us to compare genomic phenotypes across experimental conditions or clinically-defined categories. This task is made challenging by the fact that, for most genomic phenotypes, we cannot pre-specify all relevant functional units and test each unit for an overall difference across groups. For instance, for epigenetic features, many potentially relevant regulatory regions are not annotated. Furthermore, in exploratory studies, it is not always known which classes of genetic elements, such as promoters or enhancers, should be considered.

When pre-specified functional units are not available, a simple option for comparing genomic phenotypes across conditions is to perform a separate test at each locus. The majority of existing methods take such an approach, for instance using logistic regression, Fisher's exact test, or t-tests (Akalin *and others*, 2012; Stockwell *and others*, 2014; Jaffe *and others*, 2012; Hebestreit *and others*, 2013), sometimes followed by a multiplicity correction that considers spatial structure. These procedures identify differential regions by merging contiguous sites with large test statistics. Such tests tend to be under-powered, as they fail to borrow strength across neighboring loci.

In contrast, the BSmooth method of Hansen *and others* (2012) and the WaveQTL method of Shim and Stephens (2015) are two-step procedures that leverage the spatial structure of the

genomic phenotypes. BSmooth first smooths the data, and then uses the smoothed data to calculate a t -statistic at each site. Differential regions are then identified by merging contiguous sites with large t -statistics. WaveQTL requires the genome to be divided into pre-specified bins. A hierarchical Bayesian regression is performed in order to generate a bin-level test statistic, as well as estimates of association between the data and the outcome at different spatial scales.

In this paper, we propose *joint adaptive differential estimation* (JADE), a one-step approach for differential estimation and testing of genomic phenotypes. JADE is a penalized likelihood-based approach which simultaneously estimates smooth average-group profiles and identifies regions of difference between groups. By combining these two tasks into a single step, JADE can adaptively share information both across loci and between groups, leading to improved power to detect differential regions without the need for pre-specified functional units of interest. When the grouping variable has more than two levels, JADE finds regions where at least one group differs from the rest, and within those differential regions performs local clustering of profiles.

The rest of this paper is organized as follows. In Section 2, we introduce the underlying model, and formulate JADE as the solution to a convex optimization problem. In Section 3, we introduce a custom algorithm that can be used to efficiently solve the JADE optimization problem. In Section 4, we explore the performance of JADE, relative to existing methods, in a simulation study. In Section 5 we apply JADE to publicly available methylation data from the ENCODE project. The Discussion is in Section 6.

2. PROBLEM FORMULATION

Consider a categorical trait, $X \in \{1, \dots, M\}$, such as disease status or tissue type, coded numerically for convenience. We wish to associate this trait with a genomic phenotype, $\mathbf{Y} = (Y_1, \dots, Y_p)^T$, measured at positions $s_1 < s_2 < \dots < s_p$ along the genome.

For a given value of X , we assume that \mathbf{Y} varies smoothly as a function of genomic position,

$$f_m(s_j) = E[Y_j | X = m].$$

Here the function f_m represents the mean genomic phenotype profile for the m th class. If all m profiles are identical at site s_j ($f_m(s_j) = f_{m'}(s_j)$ for all $m \neq m'$), then there is no association between the mean of Y_j , the genomic phenotype at the j th position, and the categorical trait X . If $f_m(s_j) \neq f_{m'}(s_j)$ for some $1 \leq m < m' \leq M$, then there is an association between the mean of Y_j and X . Our goal is to identify *differential regions*, or contiguous blocks of associated sites. A very similar framework was considered in Shim and Stephens (2015).

In what follows, we assume that we have n independent observations of (X, \mathbf{Y}) , denoted $(x_1, \mathbf{y}_1), \dots, (x_n, \mathbf{y}_n)$. We now introduce some notation that will be used throughout this paper. Let N_m denote the number of observations with $x_i = m$, so that $N_1 + \dots + N_M = n$. Let $\bar{y}_{mj} \equiv \sum_{i: x_i = m} y_{ij} / N_m$, and let $\bar{\mathbf{y}}_m \equiv (\bar{y}_{m1}, \dots, \bar{y}_{mp})^\top$. Furthermore, we let $\theta_{mj} \equiv f_m(s_j)$, and $\boldsymbol{\theta}_m \equiv (\theta_{m1}, \dots, \theta_{mp})^\top$. In what follows, unless otherwise specified, the letter i will index the n observations, m will index the M values of the categorical trait X , and j will index the p genomic positions of Y .

2.1 Example

We illustrate JADE with a simple toy example. In each of two groups, we simulate a quantitative genomic phenotype at a series of evenly spaced positions, s_1, \dots, s_p . The data are generated as an overall group-specific mean curve, plus independent normal errors, as shown in Figure 1a. The two group-specific mean curves differ only for $s_j \in [55, 85]$.

We first consider estimating the mean curves by separately smoothing the data corresponding to each of the two groups. As is shown in Figure 1b, the two estimated profiles are somewhat different at nearly every location.

In contrast, the results from applying JADE to this data are shown in Figure 1c. JADE

simultaneously smooths the data in each group, and penalizes the differences between the two estimated mean curves. Therefore, JADE can approximately recover the differential region shown in Figure 1a.

Of course, the data that we encounter in real biological problems, such as the application studied in Section 5, are more complicated than the toy example shown in Figure 1a. Real data are often characterized by unevenly spaced positions s_1, \dots, s_p ; sites for which a subset of groups are missing measurements; and non-constant variance of the genomic phenotype measurements. As we describe in the following sections, JADE is able to accommodate all of these characteristics.

2.2 Penalties to Induce Structure

JADE combines two tasks: (i) estimation of a smooth mean curve within each group; and (ii) fusion of the mean curves across groups. Here we use the term *fusion* to describe JADE’s ability to provide mean curve estimates that are identical across multiple groups at a particular genomic position. That is, if our estimates of $f_m(s_j)$ and $f_{m'}(s_j)$ are identical for some $m \neq m'$, then we say that the estimated mean curves for the m th and m' th classes are *fused* at position s_j .

We briefly discuss the application of existing penalized regression methods to the two aforementioned tasks.

2.2.1 Smoothing a Genomic Phenotype Consider the task of smoothing a single observation of a genomic phenotype, $\mathbf{y}_i \in \mathbb{R}^p$, measured at (potentially unevenly spaced) positions $s_1 < \dots < s_p$. Given weights a_1, \dots, a_p , we consider the optimization problem

$$\underset{f}{\text{minimize}} \left\{ \frac{1}{2} \sum_{j=1}^p a_j (y_{ij} - f(s_j))^2 + \lambda P(f) \right\}. \quad (2.1)$$

The smoothed estimate, \hat{f} , minimizes the sum of two terms: a goodness-of-fit term between y_{ij} and $f(s_j)$, and a penalty term that discourages a rough or complex f . The penalty parameter λ controls the relative importance of these two terms. There are a number of options for $P(\cdot)$, such

as a smoothing spline penalty (Reinsch, 1971) or an ℓ_1 trend filtering penalty (Kim *and others*, 2009; Tibshirani, 2014).

Trend filtering induces piecewise polynomial estimates for \hat{f} , of pre-specified order k , with adaptively chosen knots. The choice of k is guided by the characteristics of the data at hand: for instance, trend filtering with $k = 0$ (also known as the fused lasso; see Tibshirani *and others*, 2005) is appropriate for the piecewise constant structure of copy number data (Tibshirani and Wang, 2008); while $k = 2$ is appropriate for relatively smooth DNA methylation data. Trend filtering is *locally adaptive*, in the sense that it can be used to fit a curve that is very smooth in one region of the domain and very rough in another; this is discussed extensively in Tibshirani (2014). This property is very attractive within the context of analyzing messy, heterogeneous, and heteroskedastic biological data. Consequently, in what follows, we take $P(\cdot)$ in (2.1) to be a trend filtering penalty.

For convenience, we now switch to using vector notation. The ℓ_1 trend filtering estimate, $\hat{\boldsymbol{\theta}}$, is the solution to the optimization problem

$$\underset{\boldsymbol{\theta} \in \mathbb{R}^p}{\text{minimize}} \left\{ \frac{1}{2} \sum_{j=1}^p a_j (y_{ij} - \theta_j)^2 + \lambda \|\mathbf{D}^{k+1,s} \boldsymbol{\theta}\|_1 \right\}, \quad (2.2)$$

where $\mathbf{D}^{k+1,s}$ is the $(p-k-1) \times p$ discrete $(k+1)$ th derivative matrix, the entries of which depend on both k and the spacing of s_1, \dots, s_p . The specific form of this matrix is detailed further in Section A of the supplementary material [SM] available at *Biostatistics* online, and in Tibshirani (2014).

The weights a_1, \dots, a_p in (2.1) and (2.2) can account for heterogeneity in the variance of y_{ij} . Setting $a_1 = \dots = a_p$ gives equal weight to each position. Setting a_j proportional to the inverse of an estimate of the variance of y_{ij} gives less weight to positions with lower quality data.

2.2.2 Fusing Genomic Phenotypes Now consider the task of fusing n observations of a genomic phenotype, $\mathbf{y}_1, \dots, \mathbf{y}_n \in \mathbb{R}^p$ — that is, we seek to encourage the estimated means to be identi-

cal at a given site. The *convex clustering* estimates, $\hat{\boldsymbol{\theta}}_1, \dots, \hat{\boldsymbol{\theta}}_n$, solve the optimization problem (Pelckmans *and others*, 2005; Hocking *and others*, 2011; Heinzl and Tutz, 2014)

$$\underset{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n \in \mathbb{R}^p}{\text{minimize}} \left\{ \sum_{i=1}^n \frac{1}{2} \sum_{j=1}^p a_{ij} (y_{ij} - \theta_{ij})^2 + \gamma \sum_{i < i'} \|\boldsymbol{\theta}_i - \boldsymbol{\theta}_{i'}\|_q \right\}. \quad (2.3)$$

In (2.3), a_{ij} is a weight for the j th locus in the i th observation.

For γ sufficiently large, the ℓ_q penalty in (2.3) will encourage similarity between $\hat{\boldsymbol{\theta}}_i$ and $\hat{\boldsymbol{\theta}}_{i'}$. In particular, if $q = 2$, then when γ is large, the entire vectors $\hat{\boldsymbol{\theta}}_i$ and $\hat{\boldsymbol{\theta}}_{i'}$ will tend to be identical, or completely fused at all sites. In this case, the set of observations for which $\hat{\boldsymbol{\theta}}_i$ are identical can be interpreted as clusters. In contrast, if $q = 1$, then a large value of γ will encourage individual elements $\hat{\theta}_{ij}$ and $\hat{\theta}_{i'j}$ to be identical. This amounts to fusing the vectors $\hat{\boldsymbol{\theta}}_i$ and $\hat{\boldsymbol{\theta}}_{i'}$ at a subset of the sites.

2.3 Joint Smoothing and Comparison with JADE

Recall from the beginning of Section 2 the problem set-up: each observation belongs to one of M categories, N_m denotes the number of observations within the m th category, and $\bar{\mathbf{y}}_m \in \mathbb{R}^p$ denotes the mean of the observations of the genomic phenotype within the m th category.

Our goal is to estimate a mean genomic phenotype profile, $\hat{\boldsymbol{\theta}}_1, \dots, \hat{\boldsymbol{\theta}}_M \in \mathbb{R}^p$, for each of the M categories. We want each mean profile to be smooth, and for the M mean profiles to be identically equal to each other for many of the loci s_1, \dots, s_p . To do this, we combine the smoothing and fusion penalties seen in (2.2) and (2.3) into a single convex optimization problem.

The JADE estimator is defined as the solution to the convex optimization problem

$$\underset{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M \in \mathbb{R}^p}{\text{minimize}} \left\{ \sum_{m=1}^M \frac{N_m}{2} \|\mathbf{A}_m (\bar{\mathbf{y}}_m - \boldsymbol{\theta}_m)\|_2^2 + \lambda \sum_{m=1}^M \|\mathbf{D}^{k+1, s} \boldsymbol{\theta}_m\|_1 + \gamma \sum_{m < m'} \|\boldsymbol{\theta}_m - \boldsymbol{\theta}_{m'}\|_1 \right\}. \quad (2.4)$$

This minimization consists of three terms: a weighted sum of squared residuals, a sum of ℓ_1 trend filtering penalties, and a clustering penalty. When the non-negative tuning parameter λ is sufficiently large, the trend filtering penalty encourages each mean profile to be smooth. Equa-

tion (2.4) could be modified to allow each of the M groups to have its own smoothness tuning parameter, $\lambda_1, \dots, \lambda_M$. For simplicity we use a single common parameter.

When the non-negative tuning parameter γ is sufficiently large, the clustering penalty encourages many of the p sites to have exactly the same value in the m th and m' th mean profiles, for $m \neq m'$. In fact, when γ is large enough, some of the p sites will have $\hat{\theta}_{1j} = \dots = \hat{\theta}_{Mj}$; these can be interpreted as regions of the genome where the mean profile is constant across the M groups. Thus, JADE simultaneously identifies regions of the genome in which the genomic phenotype is associated with the categorical variable X , and estimates smooth average profiles for each group. It accomplishes this in an efficient way that borrows strength across nearby sites, without performing a separate test at each site in the genome. Selection of λ and γ in (2.4) is discussed in Section 3.2.

In (2.4), \mathbf{A}_m are $p \times p$ diagonal weight matrices. These can be used to account for the fact that the elements of $\bar{\mathbf{y}}_m$ may have non-constant variance across the p sites, perhaps due to varying numbers of reads across the genome. Furthermore, if no data are available for the j th site in the m th group, then the j th diagonal element of \mathbf{A}_m can be set to zero.

3. SOLVING THE JADE OPTIMIZATION PROBLEM

3.1 *An Alternating Direction Method of Multipliers Algorithm for JADE*

The JADE optimization problem (2.4) is convex, so in principle, it can be solved with general-purpose convex solvers, such as SDPT3 (Tütüncü *and others*, 2003) or SeDuMi (Sturm, 1999). However, these solvers do not scale well to genome-sized problems. Therefore, we have developed an efficient custom *alternating direction method of multipliers* (ADMM; Boyd *and others*, 2010) algorithm for solving (2.4).

Our algorithm relies on the key observation by Tibshirani (2014) that the trend filtering penalty matrix $\mathbf{D}^{k+1,s}$ can be decomposed as $\mathbf{D}^{k+1,s} = \mathbf{D}^1 \tilde{\mathbf{D}}^{k,s}$, where \mathbf{D}^1 is the $(p-k-1) \times (p-k)$

first difference operator, and $\tilde{\mathbf{D}}^{k,s}$ is a $(p-k) \times p$ scaled k th-order difference operator. Details of these two matrices are provided in Section A of the SM.

Using this decomposition, we can re-write (2.4) as

$$\begin{aligned} & \underset{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M, \boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_M, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_M}{\text{minimize}} \left\{ \sum_{m=1}^M \frac{N_m}{2} \|\mathbf{A}_m(\bar{\mathbf{y}}_m - \boldsymbol{\theta}_m)\|_2^2 + \lambda \sum_{m=1}^M \|\mathbf{D}^1 \boldsymbol{\alpha}_m\|_1 + \gamma \sum_{m < m'} \|\boldsymbol{\beta}_m - \boldsymbol{\beta}_{m'}\|_1 \right\} \\ & \text{subject to} \quad \tilde{\mathbf{D}}^{k,s} \boldsymbol{\theta}_m = \boldsymbol{\alpha}_m, \quad \boldsymbol{\theta}_m = \boldsymbol{\beta}_m, \quad m = 1, \dots, M. \end{aligned} \quad (3.5)$$

The scaled augmented Lagrangian for this problem is

$$\begin{aligned} L(\boldsymbol{\theta}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{u}) &= \sum_{m=1}^M \frac{N_m}{2} \|\mathbf{A}_m(\bar{\mathbf{y}}_m - \boldsymbol{\theta}_m)\|_2^2 + \lambda \sum_{m=1}^M \|\mathbf{D}^1 \boldsymbol{\alpha}_m\|_1 + \gamma \sum_{m < m'} \|\boldsymbol{\beta}_m - \boldsymbol{\beta}_{m'}\|_1 \\ &+ \frac{1}{2} \sum_{m=1}^M \rho_{\alpha m} \left\| \tilde{\mathbf{D}}^{k,s} \boldsymbol{\theta}_m - \boldsymbol{\alpha}_m + \mathbf{u}_m^{(\alpha)} \right\|_2^2 + \frac{\rho_{\beta}}{2} \sum_{m=1}^M \left\| \boldsymbol{\theta}_m - \boldsymbol{\beta}_m + \mathbf{u}_m^{(\beta)} \right\|_2^2, \end{aligned} \quad (3.6)$$

where $\boldsymbol{\theta} \equiv (\boldsymbol{\theta}_1^\top, \dots, \boldsymbol{\theta}_M^\top)^\top$, $\boldsymbol{\alpha} \equiv (\boldsymbol{\alpha}_1^\top, \dots, \boldsymbol{\alpha}_M^\top)^\top$, and $\boldsymbol{\beta} \equiv (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_M^\top)^\top$. In (3.6), $\mathbf{u} \equiv (\mathbf{u}_1^\top, \dots, \mathbf{u}_M^\top)^\top$ is a vector of dual variables, where $\mathbf{u}_m \equiv \left(\left(\mathbf{u}_m^{(\alpha)} \right)^\top, \left(\mathbf{u}_m^{(\beta)} \right)^\top \right)^\top$ for $\mathbf{u}_m^{(\alpha)} \in \mathbb{R}^{p-k}$ and $\mathbf{u}_m^{(\beta)} \in \mathbb{R}^p$. The dual variables are broken into multiple components in order to allow for different step sizes, $\rho_{\alpha 1}, \dots, \rho_{\alpha M}$ and ρ_{β} , as this leads to faster convergence. In our implementation, we adjust the step sizes adaptively; details are in Section B.2 of the SM.

The ADMM algorithm corresponding to the scaled augmented Lagrangian (3.6) is given in Algorithm 1. The initialization in Step 1 simply amounts to solving a separate ℓ_1 trend filtering problem for each $\boldsymbol{\beta}_m$, $m = 1, \dots, M$. The update in Step 3(b) involves solving a fused lasso problem; this can be done using the algorithm of Johnson (2013). The update in Step 3(c) has an explicit form in the case of $M = 2$ groups (see Section B of the SM). For $M \geq 3$ groups, we make use of the solution of Hocking *and others* (2011).

If the output of Algorithm 1 has the property that $\beta_{1j} = \dots = \beta_{Mj}$ for some j , $1 \leq j \leq p$, then we conclude that at the j th locus, the M mean genomic phenotype profiles are identical. Due to numerical issues, however, we may not observe exact equality between β_{mj} and $\beta_{m'j}$ for $m \neq m'$. Therefore, in practice, we set a threshold ε , and conclude that β_{mj} and $\beta_{m'j}$ are equal

Algorithm 1 ADMM Algorithm For Solving the JADE Optimization Problem (2.4)

1. Initialize β_1, \dots, β_M as solutions to (2.4) with $\gamma = 0$.
2. For $m = 1, \dots, M$, initialize $\mathbf{u}_m = 0$ and $\alpha_m = \tilde{\mathbf{D}}^{k,s} \beta_m$.
3. Iterate until the convergence criteria described in Section B.2 of the SM are satisfied:

- (a) For $m = 1, \dots, M$, update

$$\boldsymbol{\theta}_m \leftarrow \left(N_m \mathbf{A}_m^\top \mathbf{A}_m + \rho_{\alpha m} \left(\tilde{\mathbf{D}}^{k,s} \right)^\top \tilde{\mathbf{D}}^{k,s} + \rho_\beta \mathbf{I} \right)^{-1} \cdot \left(N_m \mathbf{A}_m^\top \mathbf{A}_m \bar{\mathbf{y}}_m + \rho_{\alpha m} \left(\tilde{\mathbf{D}}^{k,s} \right)^\top \left(\alpha_m - \mathbf{u}_m^{(\alpha)} \right) + \rho_\beta \left(\beta_m - \mathbf{u}_m^{(\beta)} \right) \right).$$

- (b) For $m = 1, \dots, M$, update

$$\alpha_m \leftarrow \operatorname{argmin}_{\alpha_m} \left\{ \frac{1}{2} \left\| \alpha_m - \left(\tilde{\mathbf{D}}^{k,s} \boldsymbol{\theta}_m + \mathbf{u}_m^{(\alpha)} \right) \right\|_2^2 + \frac{\lambda}{\rho_{\alpha m}} \left\| \mathbf{D}^1 \alpha_m \right\|_1 \right\}.$$

- (c) For $m = 1, \dots, M$, update

$$\beta_m \leftarrow \operatorname{argmin}_{\beta_m} \left\{ \sum_{m=1}^M \frac{1}{2} \left\| \beta_m - \left(\boldsymbol{\theta}_m + \mathbf{u}_m^{(\beta)} \right) \right\|_2^2 + \frac{\gamma}{\rho_\beta} \sum_{m < m'} \left\| \beta_m - \beta_{m'} \right\|_1 \right\}.$$

- (d) For $m = 1, \dots, M$, update the dual variables by setting

$$\mathbf{u}_m^{(\alpha)} \leftarrow \mathbf{u}_m^{(\alpha)} + \tilde{\mathbf{D}}^{k,s} \boldsymbol{\theta}_m - \alpha_m, \quad \mathbf{u}_m^{(\beta)} \leftarrow \mathbf{u}_m^{(\beta)} + \boldsymbol{\theta}_m - \beta_m.$$

- (e) Update the step sizes $\rho_{\alpha 1}, \dots, \rho_{\alpha M}$ and ρ_β as described in Section B.2 of the SM, and rescale the dual variables by setting

$$\mathbf{u}_m^{(\alpha)} \leftarrow \mathbf{u}_m^{(\alpha)} \cdot \rho_{\alpha m}^{old} / \rho_{\alpha m} \quad \mathbf{u}_m^{(\beta)} \leftarrow \mathbf{u}_m^{(\beta)} \cdot \rho_\beta^{old} / \rho_\beta.$$

if the absolute difference between them is below ε . The mean genomic phenotype profile for the m th group can be obtained from θ_m in the output of Algorithm 1.

In practice, it is computationally prohibitive to solve the JADE optimization problem (2.4) on genome-sized data. Therefore, we take a pragmatic approach: we segment the genome, and apply JADE to each segment in parallel. In the methylation data application presented in Section 5, there is a natural segmentation that respects the biology of the problem. Other situations might require more arbitrary segmentation. Provided that the regions to be detected by JADE are short relative to the segmentation that we impose, we expect the segmentation to have little effect on the results.

3.2 Tuning Parameter Selection

The JADE optimization problem in (2.4) involves two non-negative tuning parameters. The parameter λ controls the smoothness of the mean genomic phenotype profiles while γ controls the amount of fusion between pairs of profiles. We take a two-stage approach to select λ and γ rather than performing a grid search over all combinations of values.

In both stages, cross-validation is performed by dividing the pM data points \bar{y}_{mj} , $m \in \{1 \dots M\}$, $j \in \{1 \dots p\}$, into l folds. For a given value of m , each fold contains a data point at every l th position, and the folds are staggered so that all m data points at a single position are not in the same fold. For example, if $M = 2$, $p = 10$, and $l = 5$, then the first fold could contain $\bar{y}_{1,1}$, $\bar{y}_{1,6}$, $\bar{y}_{2,2}$, and $\bar{y}_{2,7}$.

In the first stage, we set $\gamma = \infty$ in (2.4); this amounts to combining all of the data into a single trend filtering problem. We then perform cross-validation in order to select λ .

In the second stage of tuning parameter selection, we hold λ fixed at the value selected in the first stage, and select the tuning parameter γ using cross-validation. Additional details are provided in Section C of the SM.

In both stages of cross-validation, we apply the one-standard-error rule, selecting the largest tuning parameter value that has cross-validation error within one standard deviation of the minimum (Hastie *and others*, 2009).

4. SIMULATIONS

In Section 4.1, we consider a setting in which the genomic phenotype is continuous-valued. In Section 4.2, we consider a setting that is modeled after methylation sequence data.

4.1 Normal Simulations

4.1.1 *Simulation Set-Up* We simulate $n = 20$ observations, 10 in each of $M = 2$ groups, at $p = 300$ evenly spaced sites, s_1, \dots, s_p . (We explore other values of the sample size n in Section E.1 of the SM.) The data for the i th observation in the m th group at the j th site is generated as

$$y_{imj} = f_m(s_j) + \epsilon_{imj}, \quad (4.7)$$

where the functions f_1 and f_2 represent the mean genomic phenotype profiles for the two groups, and are displayed in Figure 2. The error terms ϵ_{imj} are generated in one of two ways:

1. *Auto-regressive model.* For $m = 1, 2$ and $i = 1, \dots, 10$,

$$z_{imj} \sim N(0, \sigma^2) \quad \text{for } j = 1 \dots 300,$$

$$\epsilon_{imj} = \begin{cases} z_{imj} & \text{if } j = 1 \\ z_{imj} + \rho z_{im(j-1)} & \text{if } j > 1 \end{cases}.$$

We consider values of $\sigma \in \{0.5, 1, 2\}$ and $\rho \in \{0, 0.2, 0.4\}$.

2. *Random effects model.* For $m = 1, 2$, $i = 1, \dots, 10$, and $j = 1, \dots, 300$,

$$b_{im} \sim N(0, \sigma_{\text{re}}^2), \quad z_{imj} \sim N(0, \sigma^2), \quad \epsilon_{imj} = b_{im} + z_{imj}. \quad (4.8)$$

In this set-up, b_{im} represents a mean shift for the i th observation in the m th group, such as one might expect as a result of a batch effect. We choose σ and σ_{re} such that $\sigma^2 + \sigma_{\text{re}}^2 = 5$,

and the proportion of variance due to random effects, $\sigma_{\text{re}}^2 / (\sigma^2 + \sigma_{\text{re}}^2)$, takes on values of 0.05, 0.1, 0.15, and 0.2.

4.1.2 Methods for Comparison In this section, we compare JADE to three t -test based methods. These methods decouple the tasks of estimating the mean genomic phenotype profiles for each of the M groups, and testing for differences between the M mean genomic phenotype profiles. These approaches assume that $M = 2$.

1. A two-sample t -statistic is calculated at each site, without first smoothing the data. This approach is used by methylKit (Akalin *and others*, 2012).
2. Each observation is smoothed using local likelihood, with the bandwidth chosen by generalized cross-validation. Then a two sample t -statistic is computed at each site, using the smoothed observations. BSmooth (Hansen *and others*, 2012) uses this strategy with a fixed bandwidth optimized for methylation data.
3. Each observation is smoothed using a quadratic smoothing spline, with the tuning parameter chosen by generalized cross-validation. Then a two sample t -statistic is computed at each site, using the smoothed observations.

The third method is included in order to understand the impact of different smoothing strategies. For all three methods, a threshold is chosen, and any site with a test statistic exceeding that threshold in absolute value is declared to have a different mean value between the M groups.

We use our own implementation in R of all three strategies, because methylKit and BSmooth are both implemented specifically for methylation count data, whereas the genomic phenotypes in this simulation study are continuous. We do not include the WaveQTL method of Shim and Stephens (2015) in our comparisons, as it requires the user to pass in pre-specified genomic regions, and does not provide a per-site assessment of the association between genomic phenotype and

category. In our application of JADE, we set the weight matrices A_1 and A_2 in (2.4) to the identity. Tuning parameters were selected according to the procedure in Section 3.2.

4.1.3 Results We now compare the performances of JADE and the three t -test-based methods described in Section 4.1.2. Before presenting these results, we briefly discuss the calculation of false and true positives for each method.

For a given value of γ in the JADE optimization problem (2.4), we declare a false positive if $\hat{\theta}_{1j} \neq \hat{\theta}_{2j}$ and $f_1(s_j) = f_2(s_j)$, and a true positive if $\hat{\theta}_{1j} = \hat{\theta}_{2j}$ and $f_1(s_j) \neq f_2(s_j)$. We fit JADE at around 100 values of γ , as described in Section C of the SM. For each value of γ considered, we calculate a true positive rate and a false positive rate.

For a given t -statistic method and a given choice of threshold, we declare a false positive if the absolute value of the t -statistic for the j th site exceeds the threshold and $f_1(s_j) = f_2(s_j)$. We declare a true positive if the absolute value of the t -statistic for the j th site exceeds the threshold and $f_1(s_j) \neq f_2(s_j)$. For each method, we calculate true positive and false positive rates for a sequence of threshold values.

Figures 3 and 4 display the average true positive rate (TPR) as a function of the false positive rate (FPR) for JADE and the three t -test-based methods, for the two error structures described in Section 4.1.1, averaged over 100 simulations. Colored points indicate the average TPR and FPR achieved with tuning parameters selected via cross-validation for JADE, or using a false discovery rate (FDR) of 10% for the t -statistic-based methods, as calculated using SLIM (Wang and others, 2011). Details of the calculation of these curves are given in Section D of the SM.

In all settings, JADE results in a higher TPR for any fixed FPR than the competing methods. We expect JADE to perform well in the random effects setting because it pools observations within each group before smoothing, thereby averaging out individual-level random effects. The JADE framework does not, however, account for the dependence between errors seen in the

auto-regressive simulations. These results show that JADE is robust, at least in this setting, to dependence between errors.

The t -statistic-based methods with an FDR cutoff of 10% tend to be more conservative than JADE with γ chosen by cross-validation: that is, they yield fewer false positives and fewer true positives. The average FPR for JADE with γ chosen by cross-validation increases for larger values of ρ in the auto-regressive settings and σ_{re} in the random effects settings.

In Section E.2 of the SM, we evaluate JADE and the three t -test methods using a different approach, in which we treat contiguous blocks of associated sites as single discoveries.

4.2 Binomial Simulations

4.2.1 Simulation Set-Up and Methods for Comparison In this section, we use methylation sequence data to motivate our simulation set-up. DNA methylation is a chemical modification that can affect cytosine residues directly followed by guanine residues (CpGs). In methylation sequencing experiments, DNA is fragmented, amplified, and bisulfite converted, a process in which non-methylated cytosines in CpGs are converted to uracil. These fragments are then sequenced, and the uracils and cytosines at each CpG are counted. Thus, at each CpG site we obtain two numbers: the number of sequenced fragments (reads) and the number of observed uracils (counts). We analyze methylation data in Section 5. In this section we consider a simple simulation mimicking the binomial character of methylation data.

As in Section 4.1.1, we simulate $n = 20$ observations, 10 in each of $M = 2$ groups at $p = 300$ evenly spaced sites. We generate the observed number of counts for the i th individual in the m th group at the j th site as

$$c_{imj} \sim \text{Binom}(n_{imj}, p_{imj}).$$

Section F of the SM describes the way in which n_{imj} , the total number of reads, is generated. No sites were permitted to have zero reads, as neither BSmooth (Hansen *and others*, 2012) nor

methylKit (Akalin *and others*, 2012) can accommodate this.

In order to generate the binomial probability p_{imj} , we first scaled and translated f_1 and f_2 , the two mean genomic phenotype profiles displayed in Figure 2, to take on values between 0 and 1. We then generated p_{imj} according to a random effects model, as follows:

$$b_{im} \sim N(0, \sigma_{\text{re}}^2), \quad p_{imj} = \begin{cases} 0 & \text{if } f_m(s_j) + b_{im} < 0 \\ 1 & \text{if } f_m(s_j) + b_{im} > 1 \\ f_m(s_j) + b_{im} & \text{otherwise} \end{cases}.$$

We consider values of $\sigma_{\text{re}} \in \{0, 0.02, 0.05, 0.07\}$.

We fit JADE using the observed proportions $y_{imj} = c_{imj}/n_{imj}$. Due to the binomial mean-variance relationship as well as the variable read depth, the variance of y_{imj} is not constant across sites or observations. We can estimate the variance of y_{imj} as

$$\hat{\sigma}_{imj}^2 = \frac{y_{imj}^*(1 - y_{imj}^*)}{n_{imj}}, \quad \text{where } y_{imj}^* = \frac{c_{imj} + 0.5}{n_{imj} + 1}. \quad (4.9)$$

Here, y_{imj}^* differs from y_{imj} in the inclusion of pseudo-counts to prevent estimates of zero variance. To accommodate these variance estimates in JADE, the diagonal elements of the matrix A_m in (2.4) were set to $1/\hat{\sigma}_{imj}$.

In what follows, we compare JADE to two existing methods for analyzing methylation data, methylKit (Akalin *and others*, 2012) and BSmooth (Hansen *and others*, 2012). These are methylation specific implementations of methods 1 and 2 in Section 4.1.2.

The local likelihood smoothing bandwidth is fixed in BSmooth, making its performance dependent on the spacing of the measurement sites. For this comparison, we used a separation between sites of five units ($s_1 = 0, s_2 = 5, s_3 = 10, \dots$). This spacing is close to what might be expected of bisulfite sequencing data in which measurements are closely spaced but not made at every base-pair.

4.2.2 *Results* Both BSmooth and methylKit produce a score for each site, quantifying the evidence that the mean profiles differ at that location. TPRs and FPRs for JADE, BSmooth, and methylKit were computed as described in Section 4.1.3. The results, averaged over 100 simulated data sets, are displayed in Figure 5. We find that JADE gives a higher TPR than BSmooth and methylKit for any fixed FPR in all settings.

5. APPLICATION TO METHYLATION DATA

In this section, we apply JADE to DNA methylation patterns during three stages of skeletal muscle cell development (myoblast, myotube, and adult muscle cells), using reduced representation bisulfite sequencing data from the ENCODE project (The Encode Project Consortium, 2012). Methylation data were described at the beginning of Section 4.2.

These cell lines have been studied extensively: in particular, ChIP-seq peaks, DNaseI peaks, and H3K27ac marks are also available. Therefore, we are able to compare the set of differentially methylated regions (DMRs) detected by JADE with previous findings, and we can assess co-localization with other functional annotations in order to validate our results. In what follows, we will make use of the fact that there is a developmental ordering to the three cell types in our data: myoblasts precede myotubes, which precede mature muscle cells.

5.1 Analysis

We compared DNA methylation in myoblasts, myotubes, and mature skeletal muscle cells. Three technical replicates from a single cell line are available for both myoblasts and myotubes, and two technical replicates are available for mature skeletal muscle. We pooled technical replicates, and set the diagonal elements of the A_m weight matrices in (2.4) equal to the inverse of the standard deviation estimates in (4.9). In this analysis, we only examined chromosome 22.

The locations at which DNA methylation can occur, CpG sites, are irregularly distributed

throughout the genome. Since there is no biological reason to smooth across very long distances containing no CpG sites, this irregular spacing provides a natural way to segment the genome. We divided the chromosome into segments such that neighboring CpG sites within a segment are separated by less than 2 kb, and the first and last CpG of each segment is measured in all three cell types. Segments with fewer than 20 CpGs were removed. This resulted in 477 segments with an average segment length of 3.0 kb and an average of 64 CpG sites per segment. Running JADE on each of the 477 segments in parallel, with 5-fold cross-validation, required computing efforts equivalent to running 120 cores for two days.

Neither methylKit nor BSmooth can be directly applied to this data, since both methods are intended for a two-group comparison, and in this data set we have three groups.

5.2 Results

5.2.1 DMRs Identified by JADE We applied JADE to each of the 477 segments on chromosome 22, with λ and γ selected using 5-fold cross-validation as described in Section 3.2 and Section C of the SM, and with $\varepsilon = 0.005$ as described in Section 3.1.

We declared a DMR as any contiguous set of CpGs at which two or more profiles are separated in the JADE output. Adjacent DMRs separated by a single CpG are combined to form a single DMR. We removed 48 regions containing 536 base-pairs (0.3% of base-pairs in DMRs) in which two pairs of profiles are fused (separation $< \varepsilon$) while the third pair is un-fused (separation $> \varepsilon$), since such a pattern does not give a valid partition of the three profiles.

JADE identified 220 DMRs in 127 segments, with an average length of 826 base-pairs. An example JADE fit is shown in Figure 6b. In this segment, two DMRs have been identified (shaded in blue). In the DMR on the left, all three profiles are separated, while on the right, the myotube and myoblast profiles are fused.

A single DMR may contain multiple partitions of the profiles, though those in Figure 6b

each contain only one partition. The 220 DMRs identified by JADE can be divided into 380 sub-regions, each of which contains only a single partition of the profiles.

5.2.2 Co-Localization of DMRs with Genetic and Epigenetic Landmarks In order to assess the quality of the DMRs detected by JADE, we evaluate their overlap with epigenetic annotations and genetic landmarks. We expect the DMRs detected by JADE to be enriched for some of these genetic features.

We consider epigenetic annotations obtained in myotubes and mature skeletal muscle, available from ENCODE. (Annotations obtained in myoblasts were not available.) These annotations include (i) transcription factor binding sites identified through ChIP-seq; (ii) active enhancer regions indicated by H3K27ac histone methylation; and (iii) DNase I hypersensitive sites, which mark open chromatin.

We also consider three types of genetic landmarks: (i) CpG islands, annotated in the UCSC genome browser. Evidence suggests that methylation in these regions affects gene expression (Bell *and others*, 2011; Illingworth and Bird, 2009). (ii) CpG island shores, defined as the 2 kb flanking regions of islands. Using BSmooth, Irizarry *and others* (2009) found that DMRs between colon cancer and healthy colon cells tend to be located in CpG island shores. (iii) The 2 kb flanking regions of transcription start sites (TSSs), annotated in the Gencode project (Harrow *and others*, 2012) and available from ENCODE annotations. These 2 kb flanking regions serve as proxies for promoter regions, which are typically located immediately upstream of the TSS.

We tested whether the number of detected DMRs overlapping each genetic feature differed from what would be expected by chance. Details are given in Section G.2 of the SM. We found that the DMRs identified by JADE are enriched in TSS flanking regions (Table 1).

Next, we restricted our analysis to the sub-regions of DMRs detected by JADE that are consistent with increasing methylation over the course of development (Myoblast \leq Myotube \leq Mature)

or decreasing methylation over the course of development (Myoblast \geq Myotube \geq Mature). These two groups account for approximately 58% of all base-pairs in detected DMRs, as discussed in Section G.1 of the SM. The results can be found in Table 1 of the SM. We found that the rate of overlap with CpG islands is higher in loss-of-methylation than in gain-of-methylation DMR sub-regions. This pattern is consistent with prior suggestions that demethylation plays a major role in up-regulating cell type specific gene expression over the course of development (Segalés *and others*, 2014; Hupkes *and others*, 2011). We also found that the rates of overlap with both DNase-I hypersensitive sites and H3K27ac modifications are higher in gain-of-methylation than in loss-of-methylation DMR sub-regions.

6. DISCUSSION

In this manuscript we propose JADE, a flexible method for the analysis of genomic phenotypes measured in two or more conditions. JADE combines smoothing and group comparison into a single optimization problem, resulting in improved power over competing methods.

In addition to gains in power for simple comparisons of two groups, JADE offers a novel approach for analyzing data with respect to a categorical outcome. Although the BSmooth and methylKit frameworks could be extended to categorical outcomes by using a one-way ANOVA, categorical outcomes are typically analyzed by performing pairwise comparisons, as in Carrió *and others* (2015), or one-versus-all comparisons. For example, tissue-specific DMRs have been identified by finding regions for which one cell type differs from the average over all other cell types (Irizarry *and others*, 2009; The Encode Project Consortium, 2012). JADE is able to identify DMRs across categorical outcomes, and determine the order and grouping of profiles within DMRs.

JADE is implemented as an R package `jadeTF`, currently available at the author's website <https://github.com/jean997/jadeTF>.

7. DESCRIPTION OF SUPPLEMENTARY MATERIALS

The reader is referred to the on-line Supplementary Materials for technical appendices.

8. FUNDING

D. W. was partially supported by a Sloan Research Fellowship, NIH Grant DP5OD009145, and NSF CAREER DMS-1252624. N.S. was supported by NIH Grant DP5OD019820. J.M. was supported by NIH Grant DP5OD019820 and NIH Grant F31HG008572. Methylation data analysis made use of resources provided by a Google Cloud Credits Award from Google Research.

REFERENCES

- AKALIN, ALTUNA, KORMAKSSON, MATTHIAS, LI, SHENG, GARRETT-BAKELMAN, FRANCINE E, FIGUEROA, MARIA E, MELNICK, ARI AND MASON, CHRISTOPHER E. (2012, oct). methylKit: a comprehensive R package for the analysis of genome-wide DNA methylation profiles. *Genome Biology* **13**(10), R87.
- BELL, JORDANA T, PAI, ATHMA A, PICKRELL, JOSEPH K, GAFFNEY, DANIEL J, PIQUE-REGI, ROGER, DEGNER, JACOB F, GILAD, YOAV AND PRITCHARD, JONATHAN K. (2011, jan). DNA methylation patterns associate with genetic and gene expression variation in HapMap cell lines. *Genome Biology* **12**(1), R10.
- BOYD, STEPHEN, PARIKH, NEAL, CHU, ERIC, PELEATO, BORJA AND ECKSTEIN, JONATHON. (2010). Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning* **3**(1), 1–122.
- CARRIÓ, ELVIRA, DÍEZ-VILLANUEVA, ANNA, LOIS, SERGI, MALLONA, IZASKUN, CASES, ILDEFONSO, FORN, MARTA, PEINADO, MIGUEL A. AND SUELVES, MÒNICA. (2015). Deconstruc-

- tion of DNA Methylation Patterns During Myogenesis Reveals Specific Epigenetic Events in the Establishment of the Skeletal Muscle Lineage. *Stem Cells* **33**(6), 2025–2036.
- HANSEN, KASPER D, LANGMEAD, BENJAMIN AND IRIZARRY, RAFAEL A. (2012, oct). BSmooth: from whole genome bisulfite sequencing reads to differentially methylated regions. *Genome Biology* **13**(10), R83.
- HARROW, J., FRANKISH, A., GONZALEZ, J. M., TAPANARI, E., DIEKHANS, M., KOKOCINSKI, F., AKEN, B. L., BARRELL, D., ZADISSA, A., SEARLE, S., BARNES, I., BIGNELL, A., BOYCHENKO, V., HUNT, T., KAY, M., MUKHERJEE, G., RAJAN, J., DESPACIO-REYES, G., SAUNDERS, G., STEWARD, C., HARTE, R., LIN, M., HOWALD, C., TANZER, A., DERRIEN, T., CHRAST, J., WALTERS, N., BALASUBRAMANIAN, S., PEI, B., TRESS, M., RODRIGUEZ, J. M., EZKURDIA, I., VAN BAREN, J., BRENT, M., HAUSSLER, D., KELLIS, M., VALENCIA, A., REYMOND, A., GERSTEIN, M., GUIGO, R. *and others.* (2012, sep). GENCODE: The reference human genome annotation for The ENCODE Project. *Genome Research* **22**(9), 1760–1774.
- HASTIE, TREVOR, TIBSHIRANI, ROBERT AND FRIEDMAN, JEROME. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Second edition. Springer.
- HEBESTREIT, KATJA, DUGAS, MARTIN AND KLEIN, HANS ULRICH. (2013, jul). Detection of significantly differentially methylated regions in targeted bisulfite sequencing data. *Bioinformatics* **29**(13), 1647–1653.
- HEINZL, FELIX AND TUTZ, GERHARD. (2014, jan). Clustering in linear-mixed models with a group fused lasso penalty. *Biometrical Journal* **56**(1), 44–68.
- HOCKING, TOBY DYLAN, JOULIN, ARMAND, BACH, FRANCIS AND VERT, JEAN-PHILIPPE. (2011). Clusterpath : An Algorithm for Clustering using Convex Fusion Penalties. *Proceedings of the 28th International Conference on Machine Learning (ICML)*.

- HUPKES, M., JONSSON, M. K. B., SCHEENEN, W. J., VAN ROTTERDAM, W., SOTOCA, A. M., VAN SOMEREN, E. P., VAN DER HEYDEN, M. A. G., VAN VEEN, T. A., VAN RAVESTEIN-VAN OS, R. I., BAUERSCHMIDT, S., PIEK, E., YPEY, D. L., VAN ZOELLEN, E. J. *and others.* (2011, nov). Epigenetics: DNA demethylation promotes skeletal myotube maturation. *The FASEB Journal* **25**(11), 3861–3872.
- ILLINGWORTH, ROBERT S. AND BIRD, ADRIAN P. (2009, jun). CpG islands - 'A rough guide'. *FEBS Letters* **583**(11), 1713–1720.
- IRIZARRY, RAFAEL A, LADD-ACOSTA, CHRISTINE, WEN, BO, WU, ZHIJIN, MONTANO, CAROLINA, ONYANGO, PATRICK, CUI, HENGMI, GABO, KEVIN, RONGIONE, MICHAEL, WEBSTER, MAREE, JI, HONG, POTASH, JAMES B, SABUNCIYAN, SARVEN *and others.* (2009, feb). The human colon cancer methylome shows similar hypo- and hypermethylation at conserved tissue-specific CpG island shores. *Nature genetics* **41**(2), 178–186.
- JAFFE, ANDREW E., MURAKAMI, PETER, LEE, HWAJIN, LEEK, JEFFREY T., FALLIN, M. DANIELE, FEINBERG, ANDREW P. AND IRIZARRY, RAFAEL A. (2012, feb). Bump hunting to identify differentially methylated regions in epigenetic epidemiology studies. *International Journal of Epidemiology* **41**(1), 200–209.
- JOHNSON, NICHOLAS A. (2013, apr). A Dynamic Programming Algorithm for the Fused Lasso and L_0 -Segmentation. *Journal of Computational and Graphical Statistics* **22**(2), 246–260.
- KIM, SEUNG-JEAN, KOH, KWANGMOO, BOYD, STEPHEN AND GORINEVSKY, DIMITRY. (2009). ℓ_1 Trend Filtering. *SIAM Review* **51**(2), 339–360.
- PELCKMANS, K., DE BRABANTER, J., SUYKENS, J. A. K. AND DE MOOR, B. (2005). Convex Clustering Shrinkage. In: *Workshop on Statistics and Optimization of Clustering Workshop (PASCAL)*, Number i.

- REINSCH, CHRISTIAN H. (1971). Smoothing by spline functions. II. *Numerische Mathematik* **16**(5), 451–454.
- SEGALÉS, JESSICA, PERDIGUERO, EUSEBIO AND MUÑOZ-CÁNOVES, PURA. (2014, sep). Epigenetic control of adult skeletal muscle stem cell functions. *FEBS Journal* **282**(9), 1571–1588.
- SHIM, HEEJUNG AND STEPHENS, MATTHEW. (2015, jul). Wavelet-based genetic association analysis of functional phenotypes arising from high-throughput sequencing assays. *The Annals of Applied Statistics* **9**(2), 665–686.
- STOCKWELL, PETER A., CHATTERJEE, ANIRUDDHA, RODGER, EUAN J. AND MORISON, IAN M. (2014, mar). DMAP: Differential methylation analysis package for RRBS and WGBS data. *Bioinformatics* **30**(13), 1814–1822.
- STURM, JOS F. (1999). Using SeDuMi 1.02, A Matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software* **11**(1-4), 625–653.
- THE ENCODE PROJECT CONSORTIUM. (2012, sep). An integrated encyclopedia of DNA elements in the human genome. *Nature* **489**(7414), 57–74.
- TIBSHIRANI, ROBERT, SAUNDERS, MICHAEL, ROSSET, SAHARON, ZHU, JI AND KNIGHT, KEITH. (2005, feb). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* **67**(1), 91–108.
- TIBSHIRANI, ROBERT AND WANG, PEI. (2008, jan). Spatial smoothing and hot spot detection for CGH data using the fused lasso. *Biostatistics* **9**(1), 18–29.
- TIBSHIRANI, RYAN J. (2014, feb). Adaptive piecewise polynomial estimation via trend filtering. *The Annals of Statistics* **42**(1), 285–323.
- TÜTÜNCÜ, R. H., TOH, K. C. AND TODD, M. J. (2003). Solving semidefinite-quadratic-linear programs using SDPT3. *Mathematical Programming, Series B* **95**(2), 189–217.

- WANG, HONG QIANG, TUOMINEN, LINDSEY K. AND TSAI, CHUNG JUI. (2011, jan). SLIM: A sliding linear model for estimating the proportion of true null hypotheses in datasets with dependence structures. *Bioinformatics* **27**(2), 225–231.

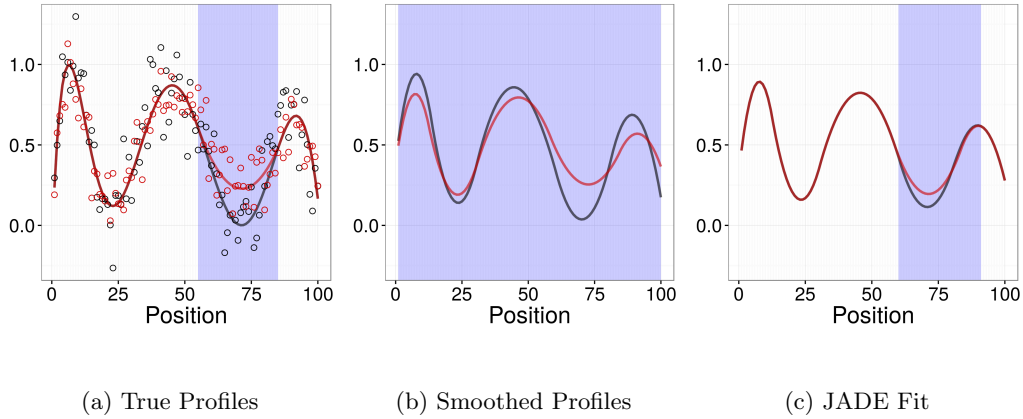


Fig. 1: An illustration of the toy example described in Section 2.1. In Figure 1a, red and black data points are generated as normal observations with mean given by the corresponding colored lines. Blue shading in 1a indicates the region in which the two true profiles are not identical. In Figure 1b, profile estimates are obtained by smoothing the two groups separately. These profiles are separated over the entire region. In Figure 1c, profile estimates are obtained from JADE. The small region in which the estimated profiles differ is shaded in blue. The detected region largely overlaps the true region of difference.

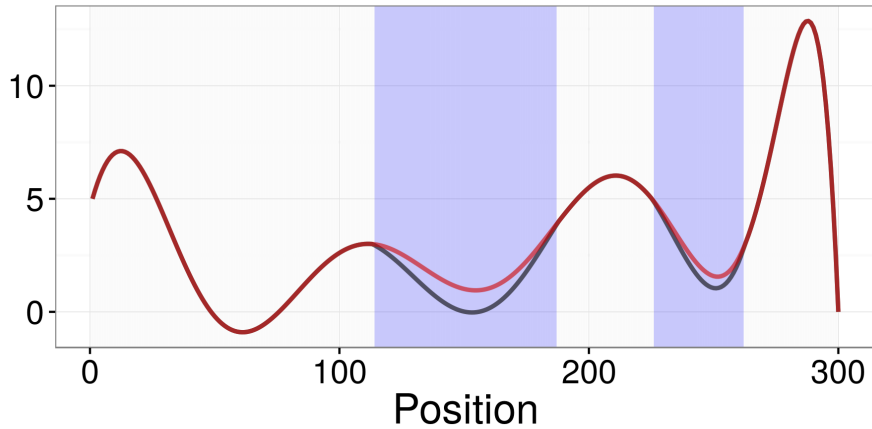


Fig. 2: Average group profiles for simulated data in Section 4. The two profiles are separated in two regions highlighted in blue. In the white regions the two groups have the same mean. For binomial data simulations in Section 4.2 these mean curves are scaled to range between 0 and 1.

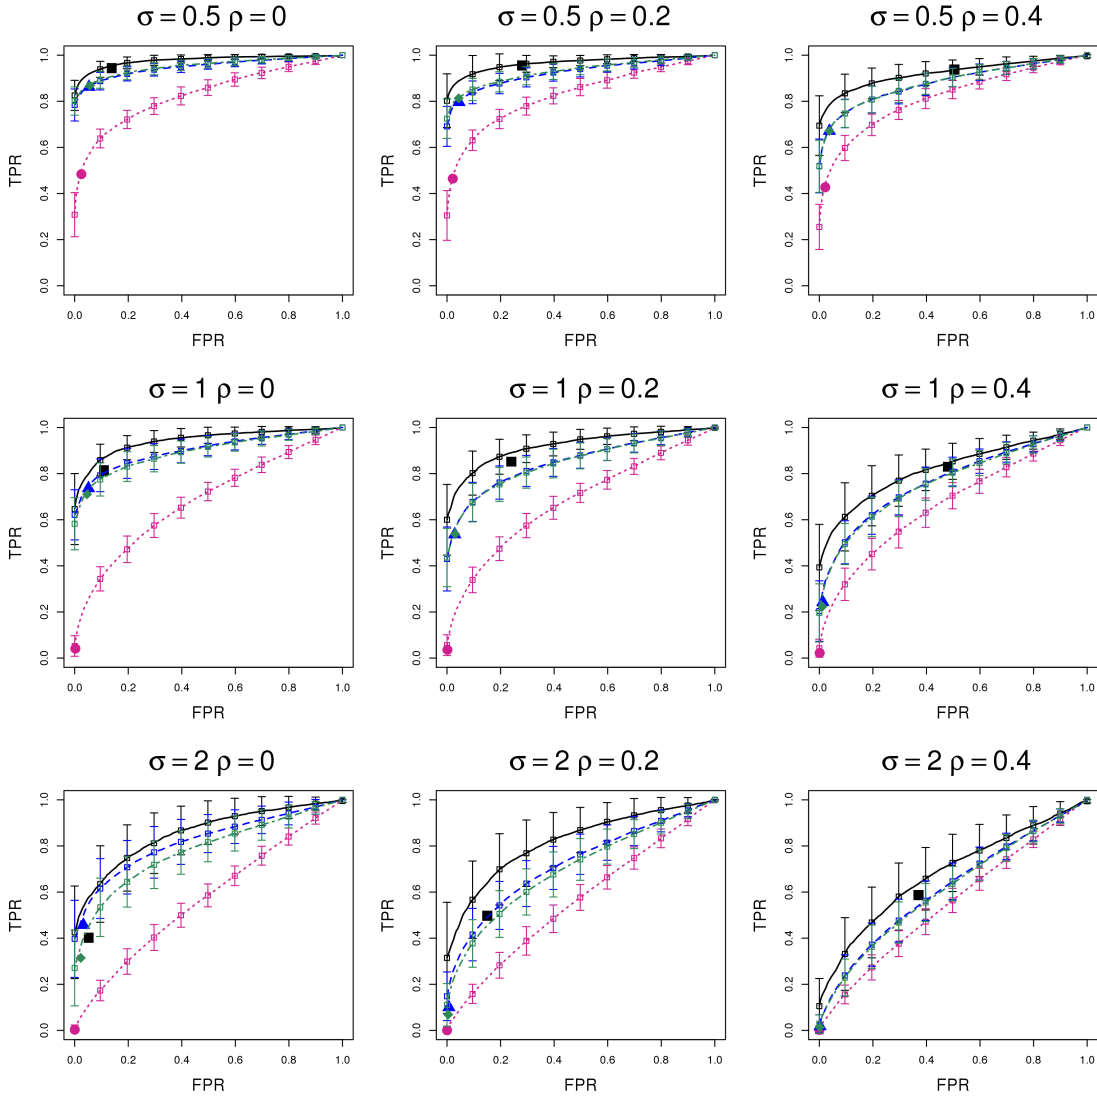


Fig. 3: Performance of JADE and competing methods in the normal auto-regressive model described in Section 4.1.1. Each panel displays results for a distinct value of $\sigma \in \{0.5, 1, 2\}$, and a value of $\rho \in \{0, 0.2, 0.4\}$. Lines show the average TPR for a fixed FPR, averaged over 100 simulations. The lengths of the vertical bars on either side of the curves equal one sample standard deviation of the TPR. Points indicate average TPR and FPR achieved for JADE with the tuning parameter selected by cross-validation, and for the t -test approaches with an FDR threshold of 10%. Methods shown are JADE (—, ■), per-site t -tests applied to the raw data (---, ●), and per-site t -tests after smoothing the raw data using splines (---, ▲) and local likelihood (---, ◆). Results for the t -test with spline and local-likelihood smoothing are often nearly identical.

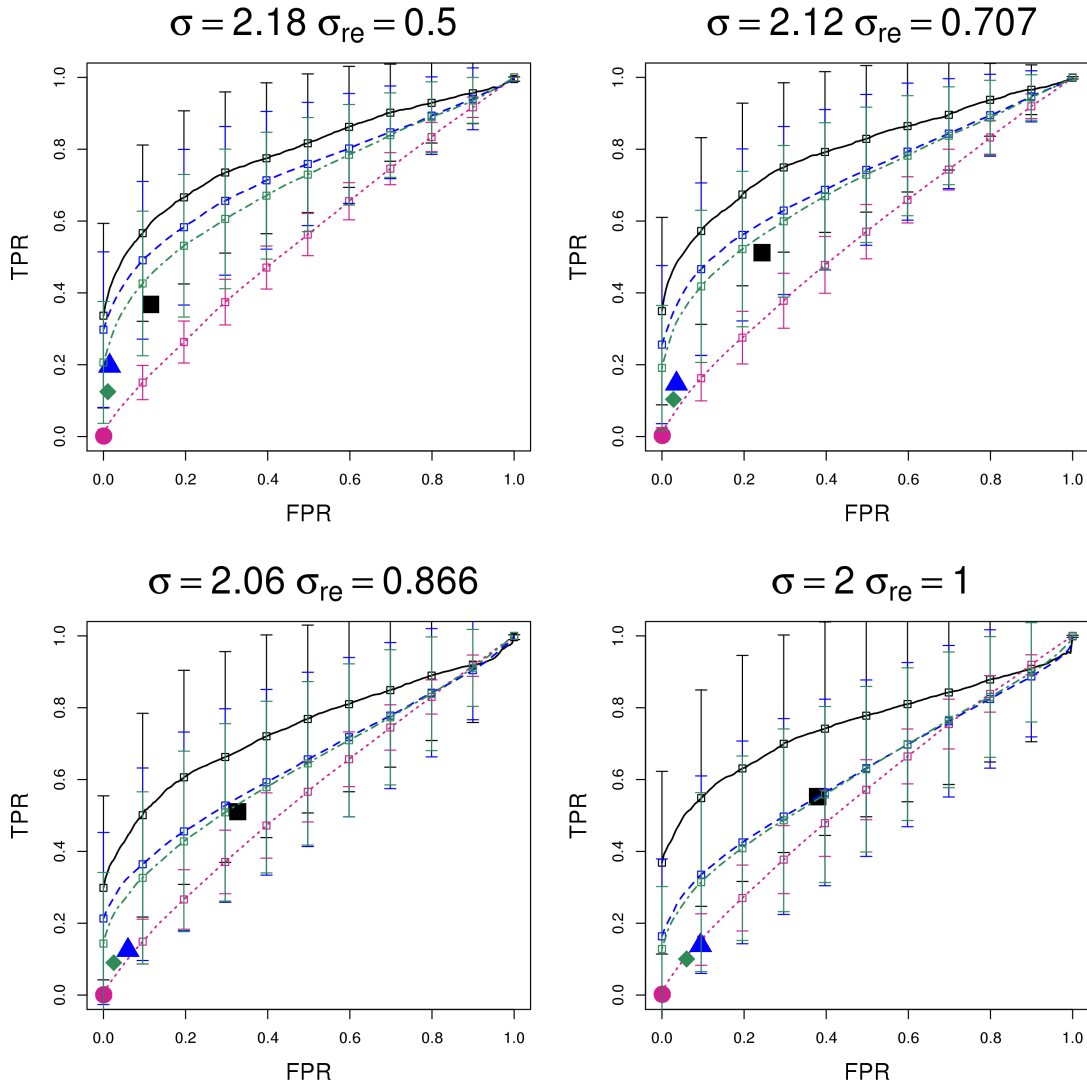


Fig. 4: Performance of JADE and competing methods in the normal random effects model described in Section 4.1.1. Each panel represents a different proportion of variation due to random effects. Additional details are as in Figure 3.

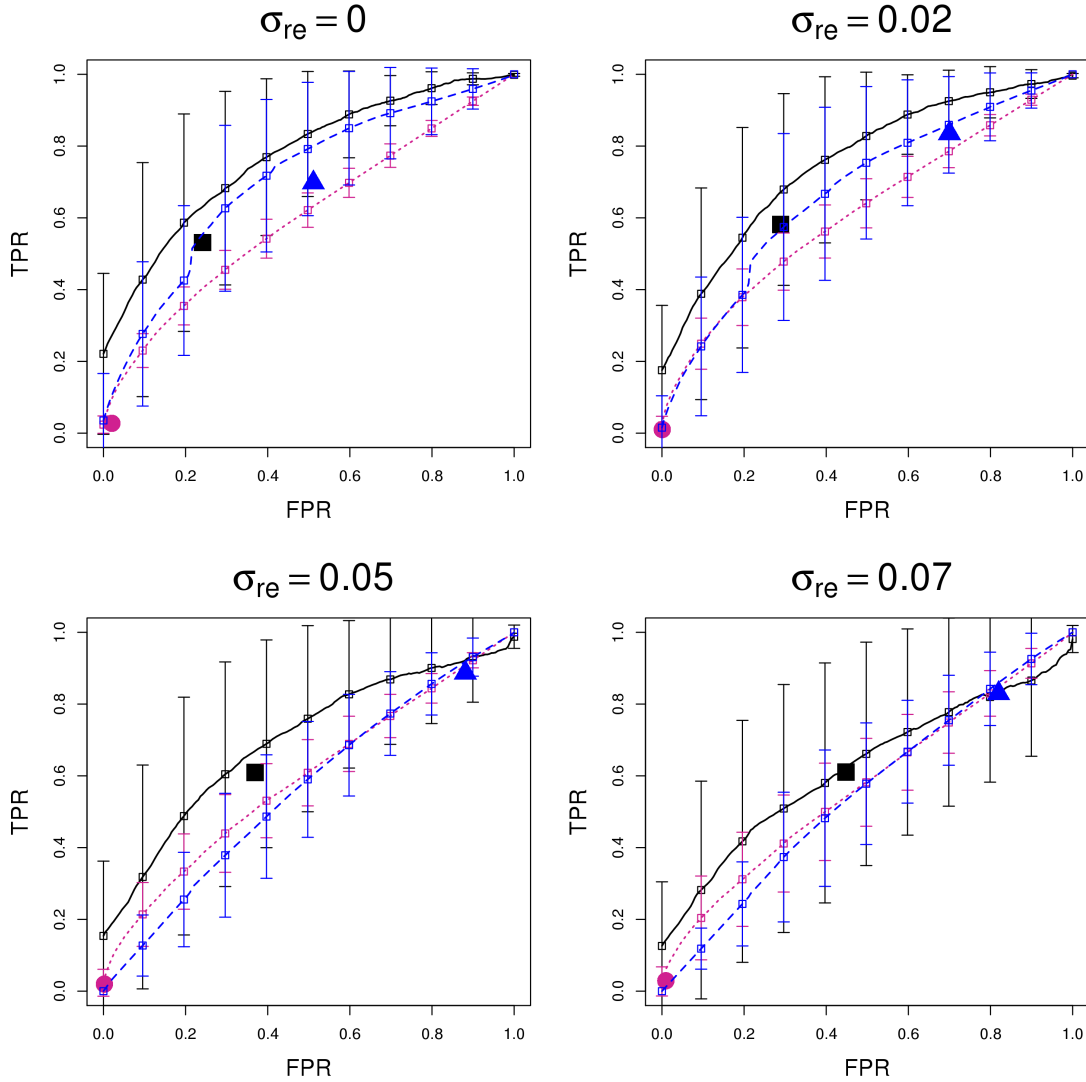


Fig. 5: Performance of JADE and competing methods in the binomial simulations described in Section 4.2.1. Lines show average TPR for a fixed FPR over 100 simulations. The lengths of the vertical bars on either side of the curves equal one sample standard deviation of the TPR. Points indicate average TPR and FPR for JADE with the tuning parameter selected by cross-validation, and for the methylKit and BSsmooth with an FDR threshold of 10%. Methods shown are JADE (—, ■), methylKit (- - -, ●), and BSsmooth (- - -, ▲).

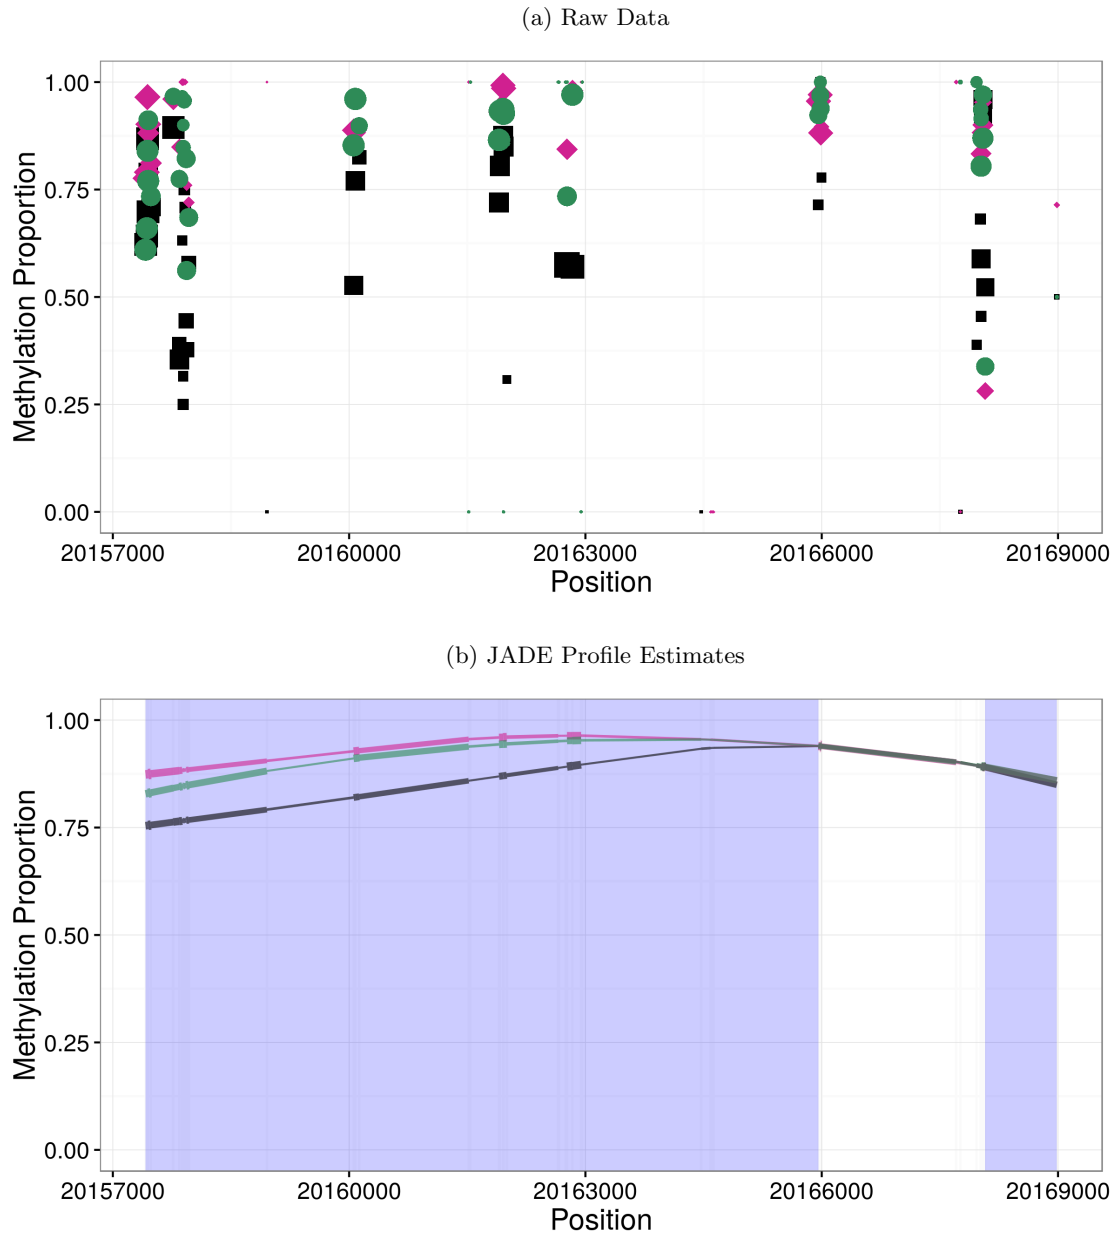


Fig. 6: Results from one segment of the methylation data analysis, described in detail in Section 5.2. Panel 6a shows the raw data for myoblasts (\blacklozenge), myotubes (\bullet) and mature skeletal muscle (\blacksquare). Point size is proportional to the number of reads at each site. Panel 6b shows the three profiles estimated by JADE (top line: myotubes, middle: myoblasts, bottom: mature skeletal muscle). Blue shading indicates DMRs detected by JADE. Line width is proportional to the number of reads in a 200 base-pair window.

Table 1: Overlap between detected DMRs and genetic features, for the methylation data analysis discussed in Section 5.2.2. ‘Total DMRs’ is the number of DMRs that overlap each genetic feature. ‘Fold’ is the ratio of the observed number of overlapping DMRs to the number that would be expected by chance. ‘P-value’ is the p-value based on a Fisher’s exact test comparing the proportion of DMRs overlapping the genetic feature to the proportion expected to occur by chance (see Section G.2 of the SM).

Genetic Feature	Total (N=220)	Fold	P-Value
Cpg Islands	119 (54.1%)	1.13	0.25
CpG Island Shores	70 (31.8%)	0.99	1
Transcription Start Sites	112 (50.9%)	1.32	0.011
TF Binding Sites	25 (11.4%)	1.03	1
DNase I HS Sites	95 (43.2%)	1.04	0.77
H3K27ac Modifications	36 (16.4%)	0.77	0.22