

# Rare Feature Selection in High Dimensions

Xiaohan Yan\*      Jacob Bien†

July 10, 2020

## Abstract

It is common in modern prediction problems for many predictor variables to be counts of rarely occurring events. This leads to design matrices in which many columns are highly sparse. The challenge posed by such “rare features” has received little attention despite its prevalence in diverse areas, ranging from natural language processing (e.g., rare words) to biology (e.g., rare species). We show, both theoretically and empirically, that not explicitly accounting for the rareness of features can greatly reduce the effectiveness of an analysis. We next propose a framework for aggregating rare features into denser features in a flexible manner that creates better predictors of the response. Our strategy leverages side information in the form of a tree that encodes feature similarity.

We apply our method to data from TripAdvisor, in which we predict the numerical rating of a hotel based on the text of the associated review. Our method achieves high accuracy by making effective use of rare words; by contrast, the lasso is unable to identify highly predictive words if they are too rare. A companion `R` package, called `rare`, implements our new estimator, using the alternating direction method of multipliers.

## 1 Introduction

The assumption of parameter sparsity plays an important simplifying role in high-dimensional statistics. However, this paper is focused on sparsity in the data itself, which actually makes estimation more challenging. In many modern prediction problems, the design matrix has many columns that are highly sparse. This arises when the features record the frequency of events (or the number of times certain properties hold). While a small number of these events may be common, there is typically a very large number of rare events, which correspond to features that are zero for nearly all observations. We call these predictors *rare features*. Rare features are in fact extremely common in many modern data sets. For example, consider the task of predicting user behavior based on past website visits: Only a small number of sites are visited by a lot of the users; all other sites are visited by only a small proportion of users.

---

\*Data Scientist, Microsoft; email: [xy257@cornell.edu](mailto:xy257@cornell.edu)

†Associate Professor, Data Sciences and Operations, Marshall School of Business, University of Southern California; email: [jbien@usc.edu](mailto:jbien@usc.edu)

As another example, consider text mining, in which one makes predictions about documents based on the terms used. A typical approach is to create a document-term matrix in which each column encodes a term’s frequency across documents. In such domains, it is often the case that the majority of the terms appear very infrequently across the documents; hence the corresponding columns in the document-term matrix are very sparse (e.g., Forman 2003; Huang 2008; Liu et al. 2010; Wang et al. 2010). In Section 6, we study a text dataset with more than 200 thousand reviews crawled from <https://www.tripadvisor.com>. Our goal is to use the adjectives in a review to predict a user’s numerical rating of a hotel. As shown in the right panel of Figure 7, the distribution of adjective density, defined as the proportion of documents containing an adjective, is extremely right-skewed, with many adjectives occurring very infrequently in the corpus. In fact, we find that more than 95% of the 7,787 adjectives appear in less than 5% of the reviews. It is common practice to simply discard rare terms,<sup>1</sup> which may mean removing most of the terms (e.g., Forman 2003; Huang 2008; Liu et al. 2010; Wang et al. 2010).

Rare features also arise in various scientific fields. For example, microbiome data measure the abundances of a large number of microbial species in a given environment. Researchers use next generation sequencing technologies, clustering these reads into “operational taxonomic units” (OTUs), which are roughly thought of as different species of microbe (e.g., Schloss et al. 2009; Caporaso et al. 2010). In practice, many OTUs are rare, and researchers often aggregate the OTUs to genus or higher levels (e.g., Zhang et al. 2012; Chen et al. 2013; Xia et al. 2013; Lin et al. 2014; Randolph et al. 2015; Shi et al. 2016; Cao et al. 2017) or with unsupervised clustering techniques (e.g. McMurdie and Holmes 2013; Wang and Zhao 2017b) to create denser features. However, even after this step, a large portion of these aggregated OTUs are still found to be too sparse and thus are discarded (e.g., Zhang et al. 2012; Chen et al. 2013; Shi et al. 2016; Wang and Zhao 2017b). The rationale for this elimination of rare OTUs is that there needs to be enough variation among samples for an OTU to be successfully estimated in a statistical model (Ridenhour et al., 2017).

The practice of discarding rare features is wasteful: a rare feature should not be interpreted as an unimportant one since it can be highly predictive of the response. For instance, using the word “ghastly” in a hotel review delivers an obvious negative sentiment, but this adjective appears very infrequently in TripAdvisor reviews. Discarding an informative word like “ghastly” simply because it is rare clearly seems inadvisable. To throw out over half of one’s features is to ignore what may be a huge amount of useful information.

Even if rare features are not explicitly discarded, many existing variable selection methods are unable to select them. The challenge is that with limited examples there is very little information to identify a rare feature as important. Theorem 1 shows that even a single rare feature can render ordinary least squares (OLS) inconsistent in the classical limit of infinite sample size and fixed dimension.

To address the challenge posed by rare features, we propose in this work a method for forming new aggregated features which are less sparse than the original ones and may be more relevant to the prediction task. Consider the following features, which represent the frequency of certain adjectives used in hotel reviews:

---

<sup>1</sup>For example, in the R text mining library `tm` (Feinerer and Hornik, 2017), `removeSparseTerms` is a commonly used function for removing any terms with sparsity level above a certain threshold.

- $\mathbf{X}_{\text{worrying}}, \mathbf{X}_{\text{depressing}}, \dots, \mathbf{X}_{\text{troubling}}$ ,
- $\mathbf{X}_{\text{horrid}}, \mathbf{X}_{\text{hideous}}, \dots, \mathbf{X}_{\text{awful}}$ .

While both sets of adjectives express negative sentiments, the first set (which might be summarized as “worry”) seems more mild than the second set (which might be summarized as “horrification”). In predicting the rating of a hotel review, we might find the following two aggregated features more relevant:

$$\begin{aligned}\widetilde{\mathbf{X}}_{\text{worry}} &= \mathbf{X}_{\text{worrying}} + \mathbf{X}_{\text{depressing}} + \dots + \mathbf{X}_{\text{troubling}} \\ \widetilde{\mathbf{X}}_{\text{horrification}} &= \mathbf{X}_{\text{horrid}} + \mathbf{X}_{\text{hideous}} + \dots + \mathbf{X}_{\text{awful}}.\end{aligned}$$

The distinction between “horrid” and “hideous” might not matter for predicting the hotel rating, whereas the distinction between a “worry”-related word versus a “horrification”-related word may be quite relevant. Thus, not only are these aggregated features less rare than the original features, but they may also be more relevant to the prediction task. A method that selects the aggregated feature  $\widetilde{\mathbf{X}}_{\text{horrification}}$  thereby can incorporate the information conveyed in the use of “hideous” into the prediction task; this same method may be unable to otherwise determine the effect of “hideous” by itself since it is too rare.

Indeed, appropriate aggregation of rare features in certain situations can be key to attaining consistent estimation and support recovery. In Theorem 2, we consider a setting where all features are rare and a natural aggregation rule exists among the features. In that setting, we show that the lasso (Tibshirani, 1996) fails to attain high-probability support recovery (for all values of its tuning parameter), whereas an oracle-aggregator does attain this property. Theorem 2 demonstrates the value of proper aggregation for accurate feature selection when features are rare. This motivates the remainder of the paper, in which we devise a strategy for determining an effective feature aggregation based on data. Our aggregation procedure makes use of side information about the features, which we find is available in many domains. In particular, we assume that a tree is available that represents the closeness of features. For example, Figure 1 shows a tree for the previous word example that is generated via hierarchical clustering over *GloVe* (Pennington et al., 2014) embeddings learned from a different data source. The two contours enclose two subtrees resulting from a cut at their joint node. Aggregating the counts in these subtrees leads to the new features  $\widetilde{\mathbf{X}}_{\text{worry}}$  and  $\widetilde{\mathbf{X}}_{\text{horrification}}$  described above. We give more details of constructing such a tree in Section 3.1.

In Section 2, we motivate our work by providing theoretical results demonstrating the difficulty that OLS and the lasso have with rare features. We further show that correct aggregation of rare features leads to signed support recovery in a setting where the lasso is unable to attain this property. In Section 3, we introduce a tree-based parametrization strategy that translates the feature aggregation problem to a sparse modeling problem. Our main proposal is an estimator formulated as a solution to a convex optimization problem for which we derive an efficient algorithm. We draw connections between this method and related approaches and then in Section 4, provide a bound on the prediction error for our method. Finally, we demonstrate the empirical merits of the proposed framework through simulation (Section 5) and through the TripAdvisor prediction task (Section 6) described above. In simulation, we examine our method’s robustness to misspecified side informa-

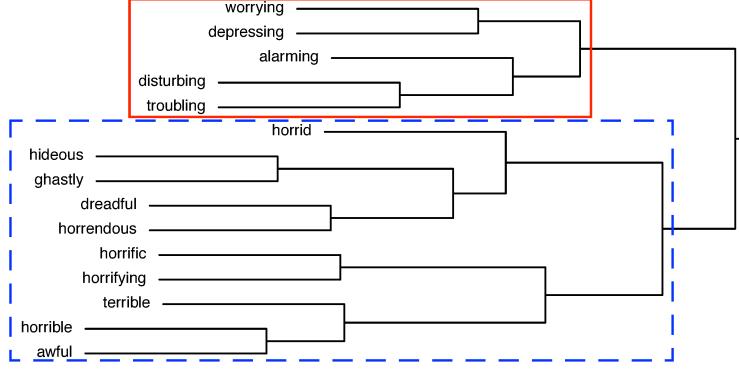


Figure 1: A tree that relates adjectives on its leaves

tion. Quantitative and qualitative comparisons in the TripAdvisor example highlight the advantages of aggregating rare features.

**Notation:** Given a design matrix  $\mathbf{X} \in \mathbb{R}^{n \times p}$ , let  $\mathbf{x}_i \in \mathbb{R}^p$  denote the feature vector of observation  $i$  and  $\mathbf{X}_j \in \mathbb{R}^n$  denote the  $j$ th feature, where  $i = 1, \dots, n$  and  $j = 1, \dots, p$ . For a vector  $\boldsymbol{\beta} \in \mathbb{R}^p$ , let  $\text{supp}(\boldsymbol{\beta}) \subseteq \{1, \dots, p\}$  denote its support (i.e., the set of indices of nonzero elements). Let  $\mathbb{S}_{\pm}(\boldsymbol{\beta}) := (\text{sign}(\beta_\ell))_{\ell=1, \dots, p}$  encode the signed support of the vector  $\boldsymbol{\beta}$ . Let  $\mathcal{T}$  be a  $p$ -leafed tree with root  $r$ , set of leaves  $\mathcal{L}(\mathcal{T}) = \{1, \dots, p\}$ , and set of nodes  $\mathcal{V}(\mathcal{T})$  of size  $|\mathcal{T}|$ . Let  $\mathcal{T}_u$  be the subtree of  $\mathcal{T}$  rooted by  $u$  for  $u \in \mathcal{V}(\mathcal{T})$ . We follow the commonly-used notions of *child*, *parent*, *sibling*, *descendant*, and *ancestor* to describe the relationships between nodes of a tree. For a matrix  $\mathbf{A} \in \mathbb{R}^{m \times n}$  and subset  $B$  of  $\{1, \dots, n\}$ , let  $\mathbf{A}_B \in \mathbb{R}^{m \times |B|}$  be the submatrix formed by removing the columns of  $\mathbf{A}$  not in  $B$ . Let  $S(\beta_\ell, \lambda) = \text{sign}(\beta_\ell) \cdot \max\{|\beta_\ell| - \lambda, 0\}$  be the soft-thresholding operator applied to  $\beta_\ell \in \mathbb{R}$ . We let  $\mathbf{e}_j$  denote the vector having a one in the  $j$ th entry and zero elsewhere. For sequences  $a_n$  and  $b_n$ , we write  $a_n \lesssim b_n$  to mean that  $a_n/b_n$  is eventually bounded, and we write  $a_n = o(b_n)$  to mean that  $a_n/b_n$  converges to zero.

## 2 Rare Features and the Promise of Aggregation

### 2.1 The Difficulty Posed by Rare Features

Consider the linear model,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I}_n). \quad (1)$$

where  $\mathbf{y} = (y_1, \dots, y_n)^T \in \mathbb{R}^n$  is a response vector,  $\mathbf{X} \in \mathbb{R}^{n \times p}$  is a design matrix,  $\boldsymbol{\beta}^*$  is a  $p$ -vector of parameters, and  $\boldsymbol{\varepsilon} \in \mathbb{R}^n$  is a vector of independent Gaussian errors having variance  $\sigma^2$ . In this paper, we focus on counts data, i.e.,  $\mathbf{X}_{ij}$  records the frequency of an event  $j$  in observation  $i$ . In particular, we will assume throughout that  $\mathbf{X}$  has non-negative elements.

The lasso (Tibshirani, 1996) is an estimator that performs variable selection, making it well-suited to the  $p \gg n$  setting:

$$\hat{\boldsymbol{\beta}}_{\lambda}^{\text{lasso}} \in \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1. \quad (2)$$

When  $\lambda = 0$ , this coincides with the OLS estimator, which is uniquely defined when  $n > p$  and  $\mathbf{X}$  is full rank:

$$\hat{\boldsymbol{\beta}}^{\text{OLS}}(n) = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}.$$

To better understand the challenge posed by rare features, we begin by considering the effect of a single rare feature on OLS in the classical  $p$ -fixed,  $n \rightarrow \infty$  regime. We take the  $j$ th feature to be a binary vector having  $k$  nonzeros, where  $k$  is a fixed value not depending on  $n$ . As  $n$  increases, the proportion of nonzero elements,  $k/n$ , goes to 0. We show in Theorem 1 that  $\hat{\beta}_j^{\text{OLS}}(n)$  does not converge in probability to  $\beta_j^*$  with increasing sample size. This establishes that OLS is not a consistent estimator of  $\boldsymbol{\beta}^*$  even in a  $p$ -fixed asymptotic regime.

**Theorem 1.** *Consider the linear model (1) with  $\mathbf{X} \in \mathbb{R}^{n \times p}$  having full column rank. Further suppose that  $\mathbf{X}_j$  is a binary vector having (a constant)  $k$  nonzeros. It follows that there exists  $\eta > 0$  for which*

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left( \left| \hat{\beta}_j^{\text{OLS}}(n) - \beta_j^* \right| > \eta \right) > 0.$$

*Proof.* The result follows from taking  $\liminf_{n \rightarrow \infty}$  of both sides of (7) in Appendix A and observing that  $2\Phi(-\eta k^{1/2}/\sigma)$  does not depend on  $n$ .  $\square$

The above result highlights the difficulty of estimating the coefficient of a rare feature. This suggests that even when rare features are not explicitly discarded, variable selection methods may fail to ever select them regardless of their strength of association with the response. Other researchers have also acknowledged the difficulty posed by rare features in different scenarios. For example, in the context of hypothesis testing for high-dimensional sparse binary regression, Mukherjee et al. (2015) shows that when the design matrix is too sparse, any test has no power asymptotically, and signals cannot be detected regardless of their strength. Since the failure is caused by the sparsity of the features, it is therefore natural to ask if “densifying the features” in an appropriate way would fix the problem. As discussed above, aggregating the counts of related events may be a reasonable way to allow a method to make use of the information in rare features.

## 2.2 Aggregating Rare Features Can Help

Given  $m$  subsets of  $\{1, \dots, p\}$ , we can form  $m$  aggregated features by summing within each subset. We can encode these subsets in a binary matrix  $\mathbf{A} \in \{0, 1\}^{p \times m}$  and form a new design matrix of aggregated features as  $\widetilde{\mathbf{X}} = \mathbf{X}\mathbf{A}$ . The columns of  $\widetilde{\mathbf{X}}$  are also counts, but represent the frequency of  $m$  different *unions* of the  $p$  original events. For example, if the first subset is  $\{1, 6, 8\}$ , the first column of  $\mathbf{A}$  would be  $\mathbf{e}_1 + \mathbf{e}_6 + \mathbf{e}_8$  and the first aggregated feature would be  $\widetilde{\mathbf{X}}_1 = \mathbf{X}_1 + \mathbf{X}_6 + \mathbf{X}_8$ , recording the number of times any of the first, sixth, or eighth events occur. A linear model,  $\widetilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}$ , based on the aggregated features can be equivalently expressed as a linear model,  $\mathbf{X}\boldsymbol{\beta}$ , in terms of the original features as long as  $\boldsymbol{\beta}$  satisfies a set of linear constraints (ensuring that it is in the column space of  $\mathbf{A}$ ):

$$\widetilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} = (\mathbf{X}\mathbf{A})\tilde{\boldsymbol{\beta}} = \mathbf{X}(\mathbf{A}\tilde{\boldsymbol{\beta}}) = \mathbf{X}\boldsymbol{\beta}.$$

The vector  $\beta$  lies in the column space of  $\mathbf{A}$  precisely when it is constant within each of the  $m$  subsets. For example,

$$\text{enforcing } \beta_1 = \beta_6 = \beta_8 \Leftrightarrow \text{aggregating features: } \mathbf{X}_1\beta_1 + \mathbf{X}_6\beta_6 + \mathbf{X}_8\beta_8 = (\mathbf{X}_1 + \mathbf{X}_6 + \mathbf{X}_8)\beta_1 = \widetilde{\mathbf{X}}_1\widetilde{\beta}_1. \quad (3)$$

In practice, determining how to aggregate features is a challenging problem, and our proposed strategy in Section 3 will use side information to guide this aggregation.

For now, to understand the potential gains achievable by aggregation, we consider an idealized case in which the correct aggregation of features is given to us by an oracle. In the next theorem, we construct a situation in which (a) the lasso on the original rare features is unable to correctly recover the support of  $\beta^*$  for any value of the tuning parameter  $\lambda$ , and (b) an oracle-aggregation of features makes it possible for the lasso to recover the support of  $\beta^*$ . For simplicity, we take  $\mathbf{X}$  as a binary matrix, which corresponds to the case in which every feature has  $n/p$  nonzero observations. We take  $\beta^*$  to have  $k$  blocks of size  $p/k$ , with entries that are constant within each block. The last block is all zeros and the minimal nonzero  $|\beta_j^*|$  is restricted to lie within a range that is expands with  $n, p$  and  $k$ . The oracle approach delivers to the lasso the  $k$  aggregated features that match the structure in  $\beta^*$ . These aggregated features have  $n/k$  nonzeros, and thus are not rare features. Having performed the lasso on these aggregated features, we then duplicate the  $k$  elements,  $p/k$  times per group, to get  $\hat{\beta}_\lambda^{\text{oracle}} \in \mathbb{R}^p$ . The lasso with the oracle-aggregator is shown to achieve high-probability signed support recovery whereas the lasso on the original features fails to achieve this property for all values of the tuning parameter  $\lambda$ .

**Theorem 2.** Consider the linear model (1) with binary matrix  $\mathbf{X} \in \{0, 1\}^{n \times p}$ ,  $p \leq n$ , and  $\mathbf{X}^T \mathbf{X} = (n/p)\mathbf{I}_p$ : every column of  $\mathbf{X}$  has  $n/p$  one's and  $\mathbf{X}\mathbf{1}_p = \mathbf{1}_n$ . Suppose  $\beta^* = \tilde{\beta}^* \otimes \mathbf{1}_{p/k}$  for  $\tilde{\beta}^* = (\tilde{\beta}_1^*, \dots, \tilde{\beta}_{k-1}^*, 0)$ . Suppose  $k < p/(36 \log n)$ . Then, the interval  $\mathcal{I} = (\sigma \sqrt{\frac{4k}{n} \log(k^2 n)}, \sigma \sqrt{\frac{p}{3n} \log(2\tilde{c}p/k)})$  with  $\tilde{c} = \frac{1}{3}e^{(\pi/2+2)^{-1}}\sqrt{\frac{1}{4} + \frac{1}{\pi}}$  is nonempty and for  $\min_{i=1, \dots, k-1} |\tilde{\beta}_i^*| \in \mathcal{I}$ , the following two statements hold:

(a) The lasso fails to get high-probability signed support recovery:

$$\limsup_{p \rightarrow \infty} \sup_{\lambda \geq 0} \mathbb{P} \left( \mathbb{S}_\pm(\hat{\beta}_\lambda^{\text{lasso}}) = \mathbb{S}_\pm(\beta^*) \right) \leq \frac{1}{e}.$$

(b) The lasso with an oracle-aggregation of features succeeds in recovering the correct signed support for some  $\lambda > 0$ :

$$\lim_{n \rightarrow \infty} \mathbb{P} \left( \mathbb{S}_\pm(\hat{\beta}_\lambda^{\text{oracle}}) = \mathbb{S}_\pm(\beta^*) \right) = 1.$$

*Proof.* See Appendix B. □

Even when the true model does not have a small number of aggregated features (i.e.,  $\beta^*$  does not have  $k \ll p$  distinct values), it may still be beneficial to aggregate. The next result exhibits a bias-variance tradeoff for feature aggregation.

**Theorem 3** (Bias-Variance Tradeoff for Feature Aggregated Least Squares). *Consider the linear model (1) with  $\mathbf{X}$  having full column rank ( $n > p$ ) and a general vector  $\boldsymbol{\beta}^* \in \mathbb{R}^p$ . Let  $\mathcal{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{C}|}\}$  be an arbitrary partition of  $\{1, \dots, p\}$ . Let  $\hat{\boldsymbol{\beta}}^{\mathcal{C}} \in \mathbb{R}^p$  be the estimator formed by performing least squares subject to the constraint that parameters are constant within the groups defined by  $\mathcal{C}$ . Then, the following mean squared estimation result holds:*

$$\frac{1}{n} \mathbb{E} \|\mathbf{X} \hat{\boldsymbol{\beta}}^{\mathcal{C}} - \mathbf{X} \boldsymbol{\beta}^*\|^2 = \frac{1}{n} \|\mathbf{X}\|_{op}^2 \sum_{\ell=1}^{|\mathcal{C}|} \sum_{j \in \mathcal{C}_\ell} \left( \beta_j^* - |\mathcal{C}_\ell|^{-1} \sum_{j' \in \mathcal{C}_\ell} \beta_{j'}^* \right)^2 + \frac{\sigma^2 |\mathcal{C}|}{n}$$

*Proof.* See Appendix C. □

The bias term is small when there is a small amount of variability in coefficient values within groups of the partition  $\mathcal{C}$ , i.e. when  $\boldsymbol{\beta}^*$  is *approximately* constant within each group. Even when  $\boldsymbol{\beta}^*$  in truth has a large  $k$  (even  $k = p$ ), there may still exist a partition  $\mathcal{C}$  with  $|\mathcal{C}| \ll k$  for which the bias term is small and thus  $\mathbb{E} \|\mathbf{X} \hat{\boldsymbol{\beta}}^{\mathcal{C}} - \mathbf{X} \boldsymbol{\beta}^*\|^2 \ll \mathbb{E} \|\mathbf{X} \hat{\boldsymbol{\beta}}^{\text{OLS}} - \mathbf{X} \boldsymbol{\beta}^*\|^2 = \sigma^2 p$ .

### 3 Main Proposal: Tree-Guided Aggregation

In the previous section, we have seen the potential gains achievable through aggregating rare features. In this section, we propose a tree-guided method for aggregating and selecting rare features. We discuss this tree in Section 3.1, introduce a tree-based parametrization strategy in Section 3.2, and propose a new estimator in Section 3.3.

#### 3.1 A Tree to Guide Aggregation

To form aggregated variables, it is infeasible to consider all possible partitions of the features  $\{1, \dots, p\}$ . Rather, we will consider a tree  $\mathcal{T}$  with leaves  $1, \dots, p$  and restrict ourselves to partitions that can be expressed as a collection of branches of  $\mathcal{T}$  (see, e.g., Figure 1). We sum features within a branch to form our new aggregated features.

We would like to aggregate features that are related, and thus we would like to have  $\mathcal{T}$  encode feature similarity information. Such information about the features comes from prior knowledge and/or data sources external to the current regression problem (i.e., not from  $y$  and  $X$ ). For example, for microbiome data,  $\mathcal{T}$  could be the phylogenetic tree encoding evolutionary relationships among the OTUs (e.g., Matsen et al. 2010; Tang et al. 2016; Wang and Zhao 2017a) or the co-occurrence of OTUs from past data sets. When features correspond to words, closeness in meaning can be used to form  $\mathcal{T}$  (e.g., in Section 6, we perform hierarchical clustering on word embeddings that were learned from an enormous corpus).

In (3), we demonstrated how aggregating a set of features is equivalent to setting these features' coefficients to be equal. To perform tree-guided aggregation, we therefore associate a coefficient  $\beta_j$  with each leaf of  $\mathcal{T}$  and “fuse” (i.e., set equal to each other) any coefficients within a branch that we wish to aggregate.

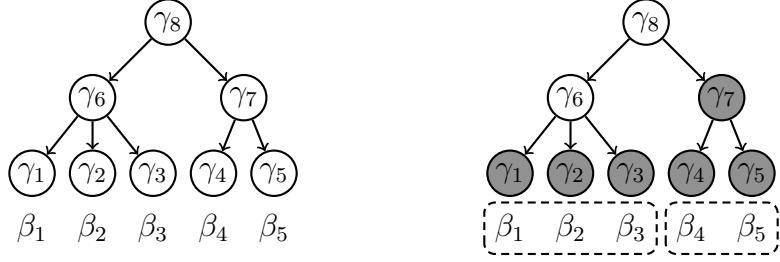


Figure 2: (Left) An example of  $\beta \in \mathbb{R}^5$  and  $\mathcal{T}$  that relates the corresponding five features. By (4), we have  $\beta_i = \gamma_i + \gamma_6 + \gamma_8$  for  $i = 1, 2, 3$  and  $\beta_j = \gamma_j + \gamma_7 + \gamma_8$  for  $j = 4, 5$ . (Right) By zeroing out the  $\gamma_i$ 's in the gray nodes, we aggregate  $\beta$  into two groups indicated by the dashed contours:  $\beta_1 = \beta_2 = \beta_3 = \gamma_6 + \gamma_8$  and  $\beta_4 = \beta_5 = \gamma_8$ . Counts data are aggregated for features sharing the same coefficient:  $\mathbf{X}\beta = (\mathbf{X}_1 + \mathbf{X}_2 + \mathbf{X}_3)\beta_1 + (\mathbf{X}_4 + \mathbf{X}_5)\beta_4$ .

### 3.2 A Tree-Based Parametrization

In order to fuse  $\beta_j$ 's within a branch, we adopt a tree-based parametrization by assigning a parameter  $\gamma_u$  to each node  $u$  in  $\mathcal{T}$  (this includes both leaves and interior nodes). The left panel of Figure 2 gives an example. Let  $ancestor(j) \cup \{j\}$  be the set of nodes in the path from the root of  $\mathcal{T}$  to the  $j$ th feature, which is associated with the  $j$ th leaf. We express  $\beta_j$  as the sum of all the  $\gamma_u$ 's on the path:

$$\beta_j = \sum_{u \in ancestor(j) \cup \{j\}} \gamma_u. \quad (4)$$

This can be written more compactly as  $\beta = \mathbf{A}\gamma$ , where  $\mathbf{A} \in \{0, 1\}^{p \times |\mathcal{T}|}$  is a binary matrix with  $A_{jk} := 1_{\{u_k \in ancestor(j) \cup \{j\}\}} = 1_{\{j \in descendant(u_k) \cup \{u_k\}\}}$ . The descendants of each node  $u$  define a branch  $\mathcal{T}_u$ , and zeroing out  $\gamma_v$ 's for all  $v \in descendant(u)$  fuses the coefficients in this branch, i.e.,  $\{\beta_j : j \in \mathcal{L}(\mathcal{T}_u)\}$ . Thus,  $\gamma_{descendant(u)} = 0$  is equivalent to aggregating the features  $\mathbf{X}_j$  with  $j \in \mathcal{L}(\mathcal{T}_u)$  (see the right panel of Figure 2).

Another way of viewing this parametrization's merging of branches is by expressing  $\mathbf{X}\beta = \mathbf{X}\mathbf{A}\gamma$ , where  $(\mathbf{X}\mathbf{A})_{ik} = \sum_{j=1}^p X_{ij} A_{jk} = \sum_{j:j \in descendant(u_k) \cup \{u_k\}} X_{ij}$  aggregates counts over all the descendant features of node  $u_k$ . By aggregating nearby features, we allow rare features to borrow strength from their neighbors, allowing us to estimate shared coefficient values that would otherwise be too difficult to estimate. In the next section, we describe an optimization problem that uses the  $\gamma$  parametrization to simultaneously perform feature aggregation and selection.

### 3.3 The Optimization Problem

Our proposed estimator  $\hat{\beta}$  is the solution to the following convex optimization problem:

$$\min_{\beta \in \mathbb{R}^p, \gamma \in \mathbb{R}^{|\mathcal{T}|}} \left\{ \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \left( \alpha \sum_{\ell=1}^{|\mathcal{T}|} w_\ell |\gamma_\ell| + (1 - \alpha) \sum_{j=1}^p \tilde{w}_j |\beta_j| \right) \quad \text{s.t. } \beta = \mathbf{A}\gamma \right\}. \quad (5)$$

We apply a weighted  $\ell_1$  penalty to induce sparsity in  $\hat{\gamma}$ , which in turn induces fusion of the coefficients in  $\hat{\beta}$ . In the high-dimensional setting, sparsity in feature coefficients is also desirable, so we also apply a weighted  $\ell_1$  penalty on  $\beta$ . The tuning parameter  $\lambda$  controls the overall penalization level while  $\alpha$  determines the trade-off between the two types of regularization: fusion and sparsity. In practice, both  $\lambda$  and  $\alpha$  are determined via cross validation. The choice of weights is left to the user. Our theoretical results (Section 4) suggest a particular choice of weights, although in our empirical studies (Sections 5 and 6) we simply take all weights to be 1 except for the root, which we leave unpenalized ( $w_{\text{root}} = 0$ ). Choosing  $w_{\text{root}} = 0$  allows one to apply strong shrinkage of all coefficients towards a common value other than zero.

When  $\alpha = 0$ , (5) reduces to a lasso problem in  $\beta$ ; when  $\alpha = 1$ , (5) reduces to a lasso problem in  $\gamma$ . Both extreme cases can be efficiently solved with a lasso solver such as `glmnet` (Friedman et al., 2010). For  $\alpha \in (0, 1)$ , (5) is a generalized lasso problem (Tibshirani and Taylor, 2011) in  $\gamma$ , and can be solved in principle using preexisting solvers (e.g., Arnold and Tibshirani 2014). However, better computational performance, in particular in high-dimensional settings, can be attained using an algorithm specially tailored to our problem. With weights  $w_r = 0$  and  $\{w_\ell = 1, \tilde{w}_j = 1\}_{\{\ell \neq r, j \in [1, p]\}}$ , we write (5) as a *global consensus problem* and solve this using alternating direction method of multipliers (ADMM, Boyd et al. (2011)). The consensus problem introduces additional copies of  $\beta$  and  $\gamma$ , which decouples the various parts of the problem, leading to efficient ADMM updates:

$$\begin{aligned} & \min_{\substack{\beta^{(1)}, \beta^{(2)}, \beta^{(3)}, \beta \in \mathbb{R}^p \\ \text{and } \gamma^{(1)}, \gamma^{(2)}, \gamma \in \mathbb{R}^{|\mathcal{T}|}}} \left\{ \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\beta^{(1)}\|_2^2 + \lambda \left( \alpha \sum_{\ell=1}^{|\mathcal{T}|} w_\ell |\gamma_\ell^{(1)}| + (1 - \alpha) \sum_{j=1}^p \tilde{w}_j |\beta_j^{(2)}| \right) \right\} \\ & \text{s.t. } \beta^{(3)} = \mathbf{A}\gamma^{(2)}, \beta = \beta^{(1)} = \beta^{(2)} = \beta^{(3)} \text{ and } \gamma = \gamma^{(1)} = \gamma^{(2)}. \end{aligned} \quad (6)$$

In particular, our ADMM approach requires performing a singular value decomposition (SVD) on  $\mathbf{X}$ , an SVD on  $(\mathbf{I}_p : -\mathbf{A})$  (these are reused for all  $\lambda$  and  $\alpha$ ), and then applying matrix multiplies and soft-thresholdings until convergence. See Algorithm 1 in Appendix D for details. Appendix D.1 provides a derivation of Algorithm 1 and Appendix D.2 discusses a slight modification for including an intercept, which is desirable in practice.

### 3.4 Connections to Other Work

Before proceeding, we draw some connections to other work. Wang and Zhao (2017b) introduce a penalized regression method with high-dimensional and compositional covariates that uses a phylogenetic tree. They apply an  $\ell_1$  penalty on the sum of coefficients within a subtree, for every possible subtree in the phylogenetic tree. Their penalty is designed to encourage the sum of coefficients within a subtree to be zero, which naturally detects a subcomposition of microbiome features. By contrast, our method applies an  $\ell_1$  penalty on our latent variables  $\gamma_u$ 's, which induces the regression coefficients within subtrees to have equal values. Thus, their penalty encourages the sum of coefficients within a subtree to be zero, whereas ours induces equality. For a simple example, suppose  $\beta_1$  and  $\beta_2$  form a subtree of the phylogenetic tree. Their penalty would promote  $\beta_1 = -\beta_2$  whereas ours would promote  $\beta_1 = \beta_2$ . The basic assumption of their method is that *contrasts* of species abundances within

a subtree may be predictive, whereas the basic assumption of our method is that *average* species abundance within a subtree may be predictive. These different structural assumptions will be appropriate in different situations. In this paper’s context of rare features, our assumption is the relevant one: Consider a subtree containing a large number of species. By asking for equality of coefficients, our method promotes the aggregation of species counts across the subtree, leading to denser features; by contrast, asking for coefficients to sum to zero does not address the problem of feature rarity.

Existing work has also considered the setting in which regression coefficients are thought to be clustered into a small number of groups of equal coefficient value. For example, She (2010) and Ke et al. (2015) do this by penalizing coefficient differences. Neither method, however, is focused specifically on rare features and thus they do not rely on side information to perform this clustering of coefficients. In our setting, the side information provided by the tree plays an important role in compensating for the extremely small amount of information available about rare features.

Several other methods assume a relevant undirected graph over the predictors and use a graph-Laplacian or graph-total-variation penalty to promote equality of coefficients that are nearby on the graph (Li and Li, 2010; Li et al., 2018). Depending on the setting, this graph may either be pure side information (Li and Li, 2010) or be a covariance graph estimated based on  $X$  itself (Li et al., 2018). While the above methods use graph information “edge-by-edge”, Yu and Liu (2016) incorporate graphical information “node-by-node” to promote joint selection of predictors that are connected on the graph.

Guinot et al. (2017) considers a similar idea of aggregating genomic features with the help of a hierarchical clustering tree; however, in the tree is learned from the design matrix and the prediction task is only used to determine the level of tree cut, whereas our method in effect uses the response to flexibly choose differing aggregation levels across the tree. We consider a strategy similar to theirs, which we call **L1-ag-h** in the empirical comparisons. Kim et al. (2012) propose a tree-guided group lasso approach in the context of multi-response regression. In their context, the tree relates the different responses and is used to borrow strength across related prediction tasks. Zhai et al. (2018) propose a variance component selection scheme that aggregates OTUs to clusters at higher phylogenetic level, and treats the aggregated taxonomic clusters as multiple random effects in a variance component model. Finally, Khabbazian et al. (2016) propose a phylogenetic lasso method to study trait evolution from comparative data and detect past changes in the expected mean trait values.

## 4 Statistical Theory

In this section, we study the prediction consistency of our method. Since  $\mathcal{T}$  encodes feature similarity information, throughout the section we require  $\mathcal{T}$  to be a “full” tree such that each node is either a leaf or possesses at least two child nodes. We begin with some definitions.

**Definition 1.** We say that  $B \subseteq \mathcal{V}(\mathcal{T})$  is an *aggregating set* with respect to  $\mathcal{T}$  if  $\{\mathcal{L}(\mathcal{T}_u) : u \in B\}$  forms a partition of  $\mathcal{L}(\mathcal{T})$ .

The black circles in Figure 3 form an aggregating set since their branches’ leaves are a partition of  $\{1, \dots, 8\}$ . We would like to refer to “the true aggregating set  $B^*$  with respect

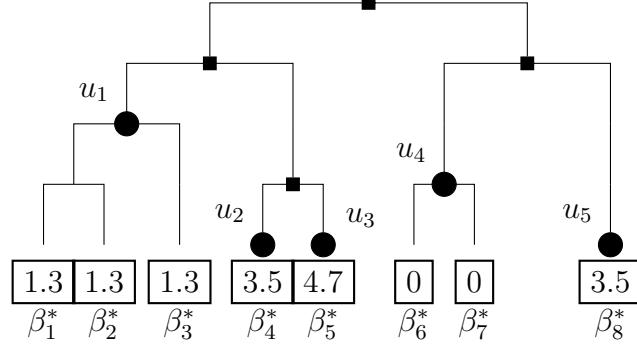


Figure 3: In the above tree,  $B^* = \{u_1, u_2, u_3, u_4, y_5\}$  has its nodes labeled with black circles.

to  $\mathcal{T}^*$  and, to do so, we must first establish that there exists a unique coarsest aggregating set corresponding to a vector  $\beta^*$ .

**Lemma 1.** *For any  $\beta^* \in \mathbb{R}^p$ , there exists a unique coarsest aggregating set  $B^* := B(\beta^*, \mathcal{T}) \subseteq \mathcal{V}(\mathcal{T})$  (hereafter “the aggregating set”) with respect to the tree  $\mathcal{T}$  such that (a)  $\beta_j^* = \beta_k^*$  for  $j, k \in \mathcal{L}(\mathcal{T}_u) \forall u \in B^*$ , (b)  $|\beta_j^* - \beta_k^*| > 0$  for  $j \in \mathcal{L}(\mathcal{T}_u)$  and  $k \in \mathcal{L}(\mathcal{T}_v)$  for siblings  $u, v \in B^*$ .*

The lemma (proved in Appendix E) defines  $B^*$  as the aggregating set such that further merging of siblings would mean that  $\beta^*$  is not constant within each subset of the partition.

**Definition 2.** Given the triplet  $(\mathcal{T}, \beta^*, \mathbf{X})$ , we define (a)  $\tilde{\mathbf{X}} = \mathbf{X} \mathbf{A}_{B^*} \in \mathbb{R}^{n \times |B^*|}$  to be the *design matrix of aggregated features*, which uses  $B^* = B(\beta^*, \mathcal{T})$  as the aggregating set, and (b)  $\tilde{\beta}^* \in \mathbb{R}^{|B^*|}$  to be the coefficient vector using these aggregated features:  $\beta^* = \mathbf{A}_{B^*} \tilde{\beta}^*$ .

We are now ready to provide a bound on the prediction error of our estimator, which is proved in Appendix F.

**Theorem 4** (Prediction Error Bound). *If we take  $\lambda \geq 4\sigma\sqrt{\log(2p)/n}$ ,  $\tilde{w}_j = \|\mathbf{X}_j\|_2/\sqrt{n}$  for  $1 \leq j \leq p$  and  $w_\ell = \|\mathbf{X}\mathbf{A}_\ell\|_2/\sqrt{n}$  for  $1 \leq \ell \leq |\mathcal{T}|$ , then*

$$\frac{1}{n} \left\| \mathbf{X} \hat{\beta} - \mathbf{X} \beta^* \right\|_2^2 \leq 3\lambda \left( (1 - \alpha) \sum_{j \in \mathcal{A}^*} \tilde{w}_j |\beta_j^*| + \alpha \sum_{\ell \in B^*} w_\ell |\tilde{\beta}_\ell^*| \right)$$

holds with probability at least  $1 - 2/p$  for any  $\alpha \in [0, 1]$ .

This is a slow rate bound for our method. The standard slow rate bound for the lasso is  $\sigma\sqrt{\log p/n} \|\beta^*\|_1$ . The next corollary establishes that our method, for any choice of  $\alpha$ , achieves this rate.

**Corollary 1.** *Suppose  $\|\mathbf{X}_j\|_2 \leq \sqrt{n}$  for  $1 \leq j \leq p$ . Then, taking  $\lambda = 4\sigma\sqrt{\log(2p)/n}$  and using the weights in Theorem 4,*

$$\frac{1}{n} \left\| \mathbf{X} (\hat{\beta} - \beta^*) \right\|_2^2 \lesssim \sigma \sqrt{\log p/n} \|\beta^*\|_1$$

holds with probability at least  $1 - 2/p$  for any  $\alpha \in [0, 1]$ .

*Proof.* See Appendix G.  $\square$

The previous corollary establishes that one does not worsen the rate by using our method over the lasso in a generic setting. But is there an advantage to using our method? Our method is designed to do well in circumstances in which  $|B^*|$  is small, that is  $\beta^*$  is mostly constant with grouping given by the provided tree. The next corollary considers the extreme case in which  $\beta^*$  is constant (and nonzero).

**Corollary 2.** *Under the conditions of Corollary 1, suppose  $\beta_1^* = \dots = \beta_p^* \neq 0$ . Taking  $\alpha \geq [1 + \|\mathbf{X}\mathbf{1}_p\|_2/(p\sqrt{n})]^{-1}$ ,*

$$\frac{1}{n} \left\| \mathbf{X}\hat{\beta} - \mathbf{X}\beta^* \right\|_2^2 \lesssim \sigma \sqrt{\frac{\log p}{n}} \|\beta^*\|_1 \cdot \frac{\|\mathbf{X}\mathbf{1}_p\|_2}{p\sqrt{n}}$$

*holds with probability at least  $1 - 2/p$ . Thus, this improves the lasso rate when  $\|\mathbf{X}\mathbf{1}_p\|_2 = o(p\sqrt{n})$ .*

*Proof.* See Appendix H.  $\square$

For certain sparse designs, the above condition holds. For example, when  $n = p$  and  $\mathbf{X} = \sqrt{n}\mathbf{I}_n$ ,  $\|\mathbf{X}\mathbf{1}_p\|_2 = \sqrt{n}\|\mathbf{1}_n\|_2 = n$  which is  $o(p\sqrt{n}) = o(n^{3/2})$ . The next proposition considers a more general sparse  $\mathbf{X}$  scenario.

**Proposition 1.** *Suppose each column of  $\mathbf{X}$  has exactly  $r$  nonzero entries chosen at random and independently of all other columns. Suppose all nonzero entries equal  $\sqrt{n/r}$  so that  $\|\mathbf{X}_j\|_2 = \sqrt{n}$  for every column  $1 \leq j \leq p$ . If  $\frac{8\log(n+1)}{3p} < \frac{r}{n} \rightarrow 0$ , then  $\|\mathbf{X}\mathbf{1}_p\|_2/(p\sqrt{n}) \rightarrow 0$  in probability.*

*Proof.* See Appendix I.  $\square$

Proposition 1 combined with Corollary 2 demonstrates that our method can improve the prediction error rate over the lasso. While this corollary focuses on the rather extreme case in which all coefficients are equal, we expect the result to be generalizable to a wider class of settings. Indeed, the empirical results of the next section suggest that our method can outperform the lasso in many settings.

## 5 Simulation Study

We start by forming a tree  $\mathcal{T}$  with  $p$  leaves. To do so, we generate  $p$  latent points in  $\mathbb{R}$  and then apply hierarchical clustering (using `hclust` in R Core Team (2016) with complete linkage). We would like the tree to partition the leaves into  $k$  clusters of varying sizes and at differing heights in the tree (which will correspond to the *true aggregation* of the features). To do so, we first generate  $k$  cluster means  $\mu_1, \dots, \mu_k \in \mathbb{R}$  with  $\mu_i = 1/i$ . The first  $k/2$  means have  $3p/(2k)$  associated latent points each, and the remaining  $k/2$  means have  $p/(2k)$  associated latent points each. The latent points associated with  $\mu_i$  are drawn independently from  $N(\mu_i, \tau^2 \min_j (\mu_i - \mu_j)^2)$ , where  $\tau = 0.05$ .

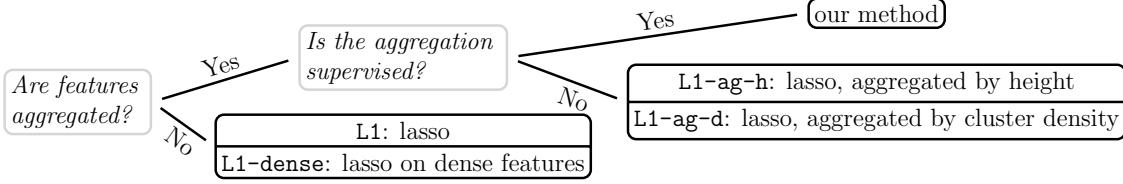


Figure 4: A comparison between our method and four other methods

By design, there are  $k$  interior nodes in  $\mathcal{T}$  corresponding to these  $k$  groups, which we index by  $B^*$ . We form  $\mathbf{A}$  corresponding to this tree and generate  $\boldsymbol{\beta}^* = \mathbf{A}_{B^*} \tilde{\boldsymbol{\beta}}^*$ . We zero out  $k \cdot s$  elements of  $\tilde{\boldsymbol{\beta}}^* \in \mathbb{R}^k$  and draw the magnitudes of the remaining elements independently from a Uniform(1.5, 2.5) distribution. We alternate signs of the nonzero coefficients of  $\tilde{\boldsymbol{\beta}}^*$ . The design matrix  $\mathbf{X} \in \mathbb{R}^{n \times p}$  is simulated from a Poisson(0.02) distribution, and the response  $\mathbf{y} \in \mathbb{R}^n$  is simulated from (1) with  $\sigma = \|\mathbf{X}\boldsymbol{\beta}^*\|_2/\sqrt{5n}$ . For every method under consideration, we average its performance over 100 repetitions in all the following simulations.

We consider both low-dimensional ( $n = 500$ ,  $p = 100$ ,  $s = 0$ ) and high-dimensional ( $n = 100$ ,  $p = 200$ ,  $s \in \{0.2, 0.6\}$ ) scenarios, in each case taking a sequence of  $k$  values up to  $p/2$ . We apply our method with  $\mathcal{T}$  and vary the tuning parameters  $(\alpha, \lambda)$  along an 8-by-50 grid of values. We take all weights equal to 1 except that of the root node, which we take as zero (leaving it unpenalized). In the low-dimensional case, we compare our method to *oracle least squares*, in which we perform least squares on  $\mathbf{X}\mathbf{A}_{B^*}$ . Oracle least squares represents the best possible performance of any method that attempts to aggregate features. We also include least squares on the original design matrix  $\mathbf{X}$ . In the high-dimensional case, we compare our method to the *oracle lasso*, in which the true aggregation  $\mathbf{X}\mathbf{A}_{B^*}$  (but not the sparsity in  $\tilde{\boldsymbol{\beta}}^*$ ) is known, and to the lasso and ridge regression, which are each computed across a grid of 50 values of the tuning parameter.

In addition to the above methods, we compare our method to three other approaches, meant to represent variations of how the lasso is typically applied when rare features are present (see Figure 4 for a schematic). The first approach, which we refer to as **L1-dense**, applies the lasso after first discarding any features that are in fewer than 1% of observations. The second and third approaches apply the lasso with features aggregated according to  $\mathcal{T}$  in an unsupervised manner. The second approach, **L1-ag-h**, aggregates features that are in the same cluster after cutting the tree at a certain height. In addition to the lasso tuning parameter, the height at which we cut the tree is a second tuning parameter (chosen along an equally-spaced grid of eight values). The third approach, **L1-ag-d**, performs merges in a bottom-up fashion along the tree until all aggregated features have density above some threshold. This threshold is an additional tuning parameter (chosen along an equally spaced grid of eight values between 0.01 and 1). The lasso tuning parameter in these methods is always chosen along a grid of 50 values.

We measure the *best mean-squared estimation error*, i.e.,  $\min_{\Lambda} \|\hat{\boldsymbol{\beta}}(\Lambda) - \boldsymbol{\beta}^*\|_2^2/p$ , where “best” is with respect to each method’s tuning parameter(s)  $\Lambda$ . The top two panels of Figure 5 show the performance of the methods in the low-dimensional and high-dimensional scenarios, respectively. Given that our method includes least squares and the lasso as special cases, it is no surprise that our methods have better attainable performance than those

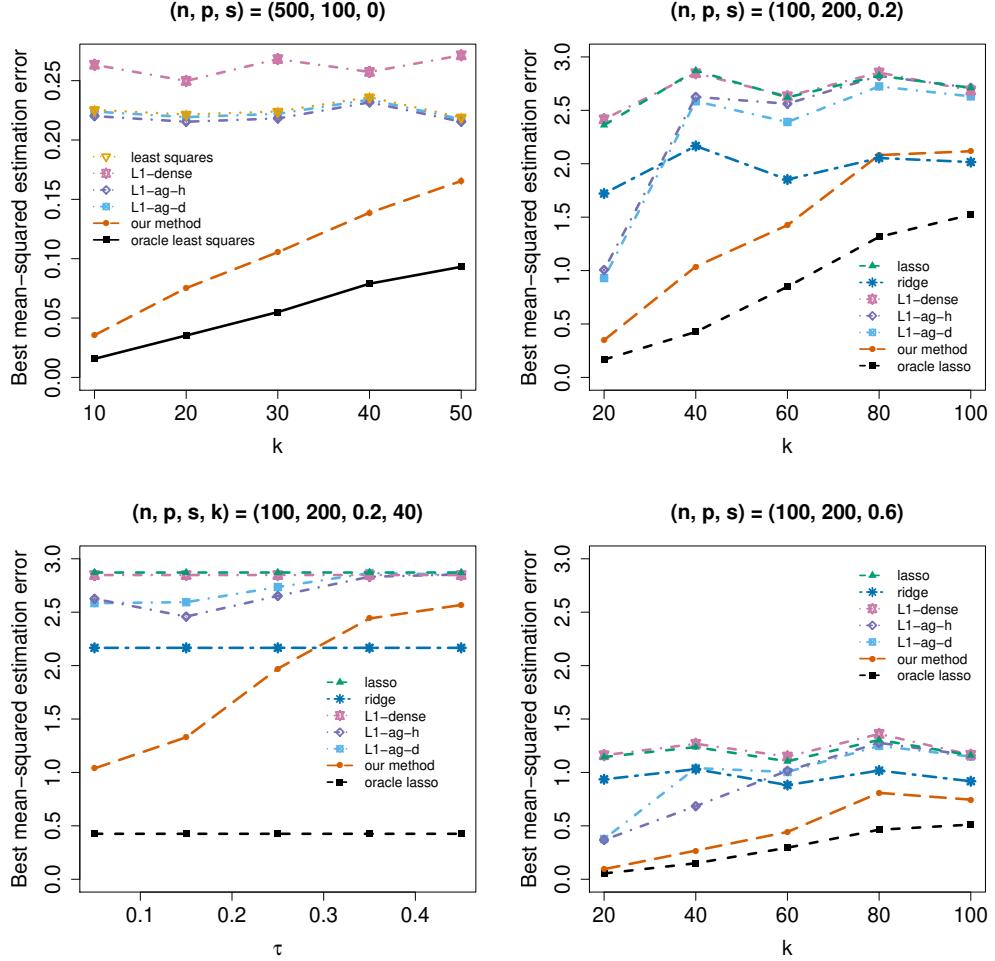


Figure 5: Estimation error of all methods under (Top Left)  $(n, p, s) = (500, 100, 0)$  versus varying  $k \in \{10, 20, 30, 40, 50\}$ ; (Top Right)  $(n, p, s) = (100, 200, 0.2)$  versus varying  $k \in \{20, 40, 60, 80, 100\}$ ; (Bottom Left)  $(n, p, s, k) = (100, 200, 0.2, 40)$  versus varying  $\tau \in \{0.05, 0.15, 0.25, 0.35, 0.45\}$ ; (Bottom Right)  $(n, p, s) = (100, 200, 0.6)$  versus varying  $k \in \{20, 40, 60, 80, 100\}$ .

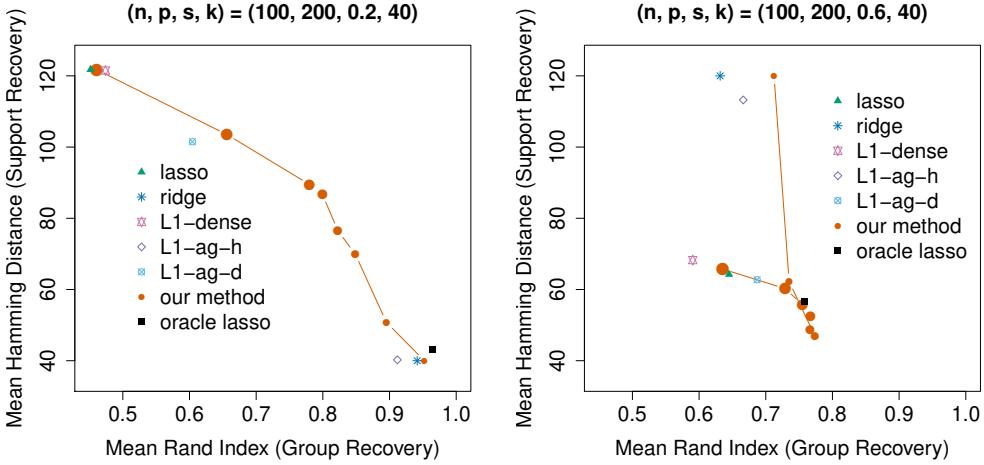


Figure 6: Mean Hamming distance versus mean Rand index for various methods' best estimates under (Left)  $(n, p, s, k) = (100, 200, 0.2, 40)$  and (Right)  $(n, p, s, k) = (100, 200, 0.6, 40)$ . For our method, the eight circles correspond to eight different  $\alpha$  values, varying from 0 to 1. Decreasing circle size corresponds to increasing  $\alpha$  value.

methods. These results indicate that our method performs nearly as well as the oracle when the true number of aggregated features,  $k$ , is small and degrades (along with the oracle) as this quantity increases. The two other methods that use the tree, L1-ag-h and L1-ag-d, do less well than our method, but still do better than L1-dense, which simply discards rare features. In the  $(n, p, s) = (100, 200, 0.2)$  case, L1-dense performs almost identically to the lasso, while L1-ag-h and L1-ag-d degrade to the lasso as  $k$  increases. By comparing the right two panels of Figure 5, we notice our method outperforms the ridge at large  $k$  when  $s$  increases from 0.2 to 0.6, which can be explained by the increased sparsity in  $\beta^*$ .

We also evaluate model performance with respect to *group recovery* and *support recovery*. Recall that our method is computed over eight  $\alpha$  values, between 0 and 1, and fifty  $\lambda$  values. At each  $\alpha$ , we find the minimizer  $\hat{\beta}(\hat{\lambda})$  for  $\|\hat{\beta}(\lambda) - \beta^*\|_2^2$  over all  $\lambda$  values. We measure group recovery by computing the Rand index (comparing the grouping of  $\hat{\beta}$  to that of  $\beta^*$ ), and measure support recovery by computing the Hamming distance (between the supports of  $\hat{\beta}$  to that of  $\beta^*$ ). The closer the Rand index is to one, the better our method recovers the correct groups. The smaller the Hamming distance is, the better our method recovers the correct support. For the high-dimensional scenario with  $k = 40$ , we plot eight pairs (one for each  $\alpha$  value) of Rand index and Hamming distance values for  $s = 0.2$  and  $s = 0.6$  in Figure 6. We also compute the two metrics for the lasso, ridge, L1-dense, L1-ag-h, L1-ag-d, and oracle lasso. In the left panel of Figure 6, which corresponds to the low-sparsity case in  $\beta^*$ , our method achieves its best performance at the largest  $\alpha$  value. As the sparsity level increases, we see from the right panel of Figure 6 that the best  $\alpha$  shifts towards zero. In both cases, our method outperforms the lasso, ridge, L1-dense, L1-ag-h, and L1-ag-d.

Clearly, the performance of our method, L1-ag-h and, L1-ag-d will depend on the quality of the tree being used. In the previous simulations we provided our method with a tree that is perfectly compatible with the true aggregating set. In practice, the tree used may be only an approximate representation of how features should be aggregated. We therefore study the

sensitivity of our method to misspecification of the tree. We return to the high-dimensional setting above with  $k = 40$ , and we generate a sequence of trees that are increasingly distorted representations of how the data should in fact be aggregated.

We begin with a true aggregation of the features into  $k$  groups as described before. In each repetition of the simulation, we generate a (random) tree  $\mathcal{T}$  by performing hierarchical clustering on  $p$  random variables generated similarly as before except having increasing  $\tau$  value, as a way to control the degradation level of the tree. When  $\tau$  is small, the latent variables will be well-separated by group so that the tree will have an aggregating set that matches the true aggregation structure (with high probability). As  $\tau \in \{0.05, 0.15, 0.25, 0.35, 0.45\}$  increases, the between-group variability becomes relatively smaller compared to within-group variability, and thus the information provided by the tree becomes increasingly weak. The bottom left panel of Figure 5 shows the degradation of our method as  $\tau$  increases when  $(n, p, s, k) = (100, 200, 0.2, 40)$ . Our method, L1-ag-h, and L1-ag-d all suffer from a poor-quality tree; the latter two degrade more quickly than ours.

## 6 Application to Hotel Reviews

Wang et al. (2010) crawled TripAdvisor.com to form a dataset<sup>2</sup> of 235,793 reviews and ratings of 1,850 hotels by users between February 14, 2009 and March 15, 2009. While there are several kinds of ratings, we focus on a user’s overall rating of the hotel (on a 1 to 5 scale), which we take as our response. We form a document-term matrix  $\mathbf{X}$  in which  $\mathbf{X}_{ij}$  is the number of times the  $i$ th review uses the  $j$ th adjective.

We begin by converting words to lower case and keeping only adjectives (as determined by WordNet Fellbaum 1998; Wallace 2007; Feinerer and Hornik 2016). After removing reviews with missing ratings, we are left with 209,987 reviews and 7,787 distinct adjectives. The left panel of Figure 7 shows the distribution of ratings in the data: nearly three quarters of all ratings are above 3 stars. The extremely right-skewed distribution in the right panel of Figure 7 shows that all but a small number of adjectives are highly rare (e.g., over 90% of adjectives are used in fewer than 0.5% of reviews).

Rather than discard this large number of rare adjectives, our method aims to make productive use of these by leveraging side information about the relationship between adjectives. We construct a tree capturing adjective similarity as follows. We start with word embeddings<sup>3</sup> in a 100-dimensional space that were pre-trained by *GloVe* (Pennington et al., 2014) on the Gigaword5 and Wikipedia2014 corpora. We also obtain a list of adjectives, which the NRC Emotion Lexicon labels as having either positive or negative sentiments (Mohammad and Turney, 2013). We use five nearest neighbors classification within the 100-dimensional space of word embeddings to assign labels to the 5,795 adjectives that have not been labeled in the NRC Emotion Lexicon. This sentiment separation determines the two main branches of the tree  $\mathcal{T}$ . Within each branch, we perform hierarchical clustering of the word embedding vectors. Figure 8 depicts such a tree with 2,397 adjectives (as leaves).

We compare our method (with weights all equal to 1 except for the root node, which is left unpenalized) to four other approaches described in Figure 4. For L1-dense, we first

<sup>2</sup>Data source: <http://times.cs.uiuc.edu/~wang296/Data/>

<sup>3</sup>Data source: <http://nlp.stanford.edu/data/glove.6B.zip>

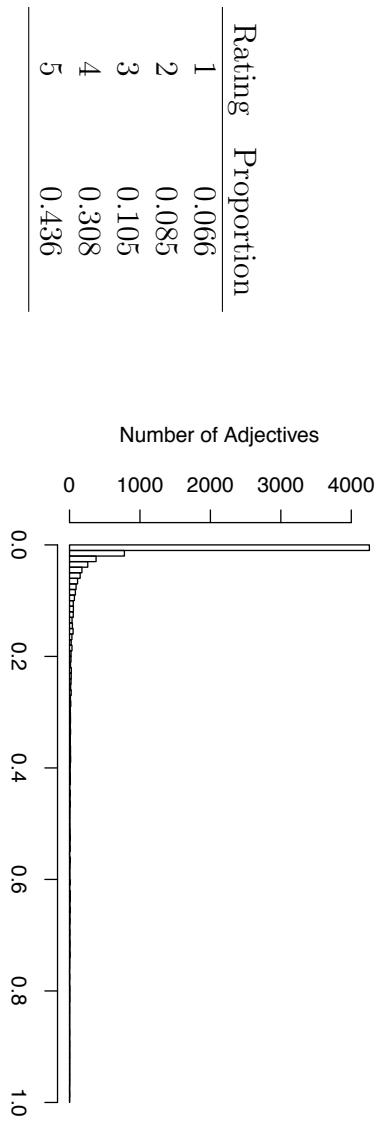


Figure 7: (Left) distribution of TripAdvisor ratings. (Right) only 414 adjectives appear in more than 1% of reviews; the histogram gives the distribution of usage-percentages for those adjectives appearing in fewer than 1% of reviews.

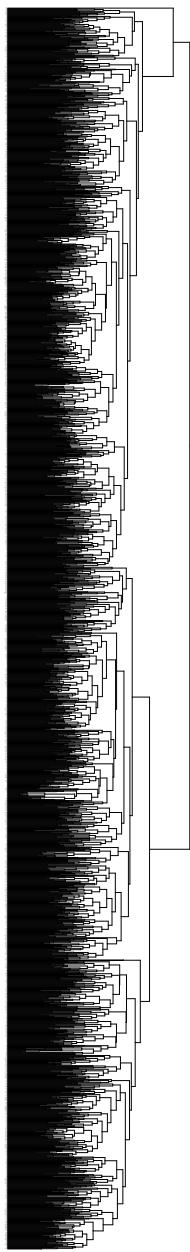


Figure 8: Tree  $\mathcal{T}$  over 2,397 adjectives: the left subtree is for adjectives with negative sentiment and the right subtree is for adjectives with positive sentiment.

discard any adjectives that are in fewer than 0.5% of reviews before applying the lasso. Both **L1-ag-h** and **L1-ag-d** have an additional tuning parameter to take care of: for **L1-ag-h** we vary the height at which we cut the tree along an equally-spaced grid of ten values; for **L1-ag-d** we choose the threshold for aggregations' density along an equally spaced grid of ten values between 0.001 and 0.1.

prop.	$n$	$p$	$n/p$	Mean Squared Prediction Error				
				our method	L1	L1-dense	L1-ag-h	L1-ag-d
1%	1,700	2,397	0.71	<b>0.870</b>	0.894	0.895	0.882	0.971
5%	8,499	3,962	2.15	<b>0.783</b>	0.790	0.805	0.785	0.899
10%	16,999	4,786	3.55	<b>0.758</b>	0.764	0.788	0.764	0.902
20%	33,997	5,621	6.05	<b>0.742</b>	0.749	0.773	0.747	1.173
40%	67,995	6,472	10.51	<b>0.739</b>	0.740	0.768	0.742	1.108
60%	101,992	6,962	14.65	<b>0.733</b>	0.736	0.769	0.734	1.155
80%	135,990	7,294	18.64	<b>0.733</b>	<b>0.733</b>	0.765	0.734	0.886
100%	169,987	7,573	22.45	<b>0.729</b>	0.731	0.765	0.731	0.956

Table 1: Performance of five methods on the held-out test set: L1 is the lasso; L1-dense is the lasso on only dense features; L1-ag-h is the lasso with features aggregated based on height; and L1-ag-d is the lasso with features aggregated based on density level.

We hold out 40,000 ratings and reviews as a test set. To observe the performance of these methods over a range of training set sizes, we consider a nested sequence of training sets, ranging from 1% to 100% of the reviews not included in the test set. For all methods, we use five-fold cross validation to select tuning parameters and threshold all predicted ratings to be within the interval [1, 5]. Table 1 displays the mean squared prediction error (MSPE) on the test set for each method and training set size.

As the size of the training set increases, all methods except for the lasso with aggregation based on density (L1-ag-d) achieve lower MSPE. Among the four lasso-related methods, L1 and L1-ag-h outperform the other two. As the training set size  $n$  increases, the number of features  $p$  also increase but at a relatively slower rate. We notice that when  $n/p$  is less than 10.51, our method outperforms the other four lasso-related methods. As  $n/p$  increases beyond 10.51, i.e., in the statistically easier regimes, L1 and L1-ag-h attain performance comparable to our method. We conduct paired  $t$ -tests between squared prediction errors from our method and L1-ag-h at every  $(n, p)$  pair (i.e., every row of Table 1). Five out of the eight tests are significant at the 0.005 significance level (See Appendix J Table 4).

To better understand the difference between our method, the lasso, and L1-ag-h, we color the branches of the tree generated in the  $n = 1,700$  and  $p = 2,397$  case (i.e., proportion is 1%) according to the sign and magnitude of  $\hat{\beta}$  for the three methods. The bottom tree in Figure 9 corresponds to our method and has many nearby branches sharing the same red/blue color, indicating that the corresponding adjective counts have been merged. By contrast, the top tree in Figure 9, which corresponds to the lasso, shows that the solution is sparser and does not have branches of similar color. The middle tree in Figure 9 shows that L1-ag-h produces something between the two, merging some adjectives with strong signals while keeping the rest as singletons. The strength and pattern of aggregations vary

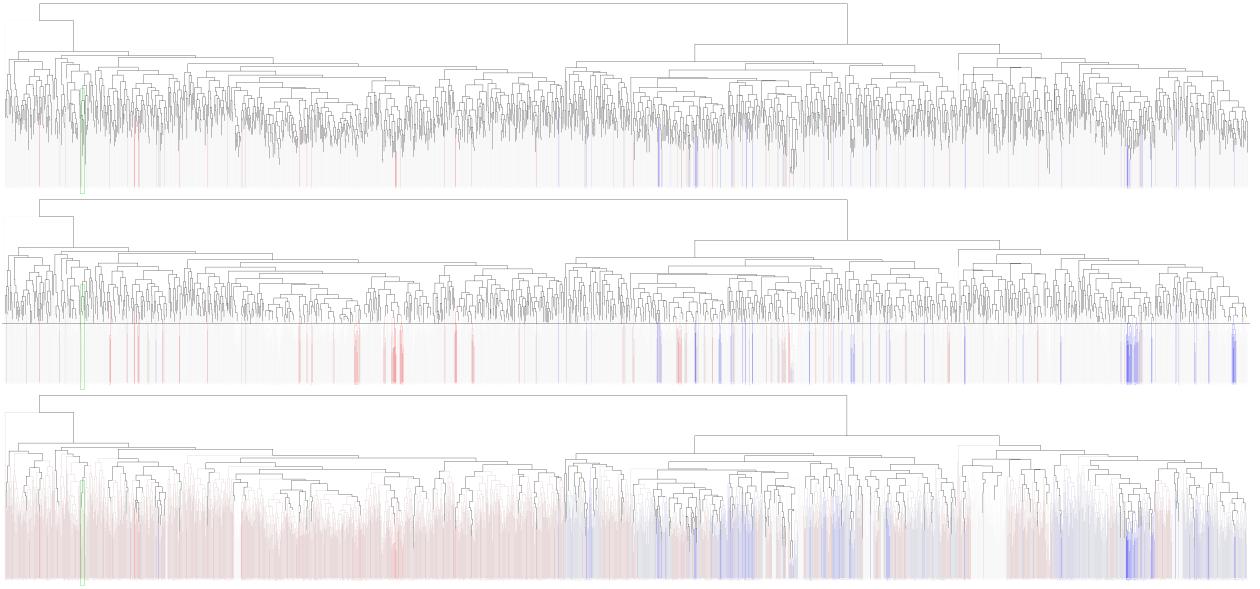


Figure 9: Trees for 2,397 adjectives on the leaves with branches colored based on  $\hat{\beta}$  estimated with the lasso (Top), L1-ag-h (Middle) and our method (Bottom), respectively. Red branch, blue branch and gray branch correspond to negative, positive and zero  $\hat{\beta}_j$ , respectively. Darker color indicates larger magnitude of  $\hat{\beta}_j$  and lighter color indicates smaller magnitude of  $\hat{\beta}_j$ . The horizontal line shown in the middle plot corresponds to the height, chosen from CV, at which L1-ag-h cuts the tree and merges the resulting branches.

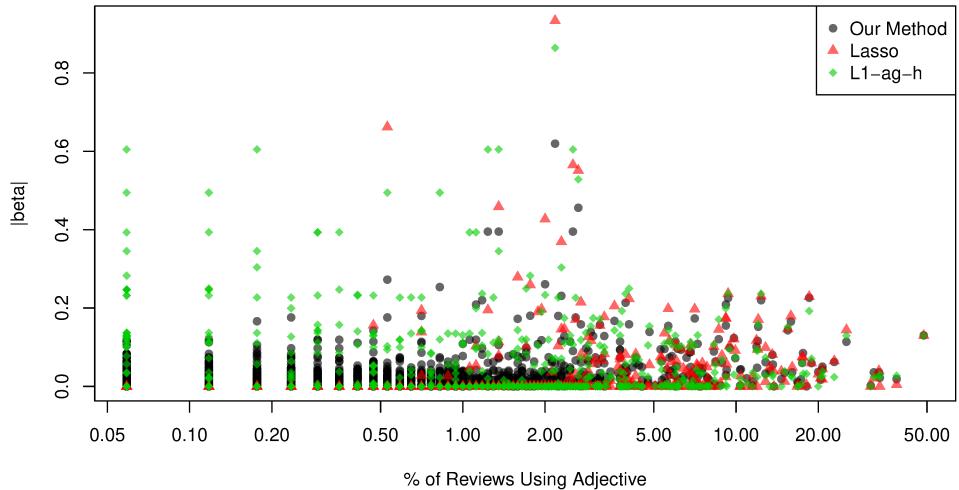


Figure 10: Magnitudes of coefficient estimates  $\{|\hat{\beta}_j|\}$  versus feature density (on log scale) for our method (black circles), the lasso (red triangles) and L1-ag-h (green diamond).

between our method and L1-ag-h. Inspection of the merged branches from our method reveals words of similar meaning and sentiment being aggregated. In Figure 10, we plot  $\{|\hat{\beta}_j|\}$  against the percentage of reviews containing an adjective. We find that our method selects rare words more than the other two methods. The rarest word selected by the lasso is “filthy”, which appears in 0.47% of reviews. By contrast, our method selects many words that are far more rare: at the extreme of rarest words, our method selects 797 words that appear in only 0.059% of reviews; L1-ag-h selects 31 out of these 797 rarest words. Our method is able to select more rare words through aggregation: it aggregates 2,244 words into 224 clusters, leaving the remaining 153 words as singletons. Over 70% of these singletons are dense words (where, for this discussion, we call a word “dense” if it appears in at least 1% of reviews and “rare” otherwise). This is four times higher than the percentage of dense words in the original training data. Of the 224 aggregated clusters, 42% are made up entirely of rare words. After aggregation, over half of these clusters become dense features. As a comparison, L1-ag-h aggregates 1,339 words into 506 clusters while keeping 1,058 words as singletons. Only 10% of these singletons from L1-ag-h are dense words. Of the 506 aggregated clusters from L1-ag-h, 68% are made by rare words only, among which only 15% become dense clusters after aggregation.

Table 2 shows the density and estimated coefficient values for eight words falling in a particular subtree of  $\mathcal{T}$ . The words “heard” and “loud” occur far more commonly than the other six words. We see that the lasso only selects these two words whereas our method selects all eight words (assigning them negative coefficient values). In contrast, L1-ag-h, while also aggregating the two densest words in this branch into a group, does not select the six rare words. Examining the six rare words, it seems quite reasonable that they should receive negative coefficient values. This suggests to us that the lasso’s exclusion of these words has to do with their rareness in the dataset rather than their irrelevance to predicting the hotel rating.

adjectives	heard	loud	yelled	shouted	screaming	crying	blaring	banging
density <sup>4</sup>	0.0300	0.0235	0.0006	0.0006	0.0029	0.0006	0.0006	0.0041
$\hat{\beta}_\lambda^{lasso}$	-0.057	-0.147	0	0	0	0	0	0
$\hat{\beta}_{\lambda, \text{height}}^{\text{L1-ag-h}}$	-0.174	-0.174	0	0	0	0	0	0
$\hat{\beta}_\lambda^{ours}$	-0.128	-0.128	-0.039	-0.039	-0.039	-0.039	-0.039	-0.039

Table 2: Term density and estimated coefficient for adjectives in the selected group

Existing work in sentiment analysis uses side information in other ways to improve predictive performance (Thelwall et al., 2010). For example, Wang et al. (2016), working in an SVM framework, forms a directed graph over words that expresses their “sentiment strength” and then requires that the coefficients corresponding to these words honor the ordering that is implied by the directed graph. For example if word A is stronger (in expressing a certain sentiment) than word B, which in turn is stronger than word C, then their method enforces  $\beta_A \geq \beta_B \geq \beta_C$ . Such constraints can in some situations have a regularizing effect that may be of use in rare feature settings: for example, if  $\beta_A \approx \beta_C$  and B is a rare word, this

<sup>4</sup>The term density is computed over the training set.

constraint would help pin down  $\beta_B$ ’s value. However, if  $\beta_A$  and  $\beta_C$  are very different from each other, the constraint may offer little help in reducing the variance of the estimate of  $\beta_B$ . Another difference is that the method uses hard constraints that are not controlled by a tuning parameter, so that even when there is strong evidence in the data that a constraint should be violated, this will not be allowed. By contrast, our method shrinks toward the constant-on-subtrees structure without forcing this to be the case.

tree setting	Mean Squared Prediction Error				
	our method	L1	L1-dense	L1-ag-h	L1-ag-d
<i>GloVe</i> -50d	<b>0.748</b>	0.749	0.773	0.752	0.775
<i>GloVe</i> -100d	<b>0.742</b>	0.749	0.773	0.747	1.173
<i>GloVe</i> -200d	<b>0.741</b>	0.749	0.773	0.747	1.140
ELMo	<b>0.669</b>	0.676	0.732	0.685	0.783

Table 3: Performance of five methods under various tree settings on the held-out test set. Among the tree settings, *GloVe*-50d, *GloVe*-100d and *GloVe*-200d correspond to hierarchical clustering trees generated with *GloVe* embeddings of differing dimensions. We collapse over two million ELMo embedding vectors of adjectives into 6,001 clusters using mini-batch K-Means clustering (Sculley, 2010), then generate a hierarchical tree upon the 6,001 cluster centroids. The performance is based on 33,997 training reviews (i.e., corresponding to the 20% row of Table 1).

Section 5 investigated the effect of using a distorted tree (see the bottom left panel of Figure 5). In the context of words there are multiple choices of trees one could use. Table 3 shows the effect of applying our method to different types of trees. Here we focus on one situation with 33,997 training reviews. We first consider the effect of changing the dimension of the *GloVe* embedding. One might suppose that as the embedding dimension increases the tree becomes more informative. Indeed, one finds that our method and L1-ag-h achieve improved performance as this dimension increases. The last method, L1-ag-d, has more variability in its performance as the *GloVe* tree varies, making it the poorest performing method. Both L1 and L1-dense do not make use of the tree and thus they are unaffected by changes to the embedding dimension. But a limitation of the methods is that they all are based on the bag-of-words model, and therefore do not take word context into consideration. For example, the models do not differentiate between the use of “bad” in “the hotel is bad” versus “the hotel is not that bad”. To bring context into consideration, we leverage deep contextualized word representations from ELMo (Peters et al., 2018) to generate an auxiliary tree (See Appendix K for details). Unlike traditional word embeddings such as *GloVe* that associate one embedding per word, deep word embeddings can capture the meaning of each word based on the surrounding context. From Table 3 we find test errors improve substantially with the ELMo tree for our method, L1 and L1-ag-h. Among all methods and all tree settings, our method with the ELMo tree performs the best. This suggests our method can better leverage the power of contextualized word embeddings than competing methods.

## 7 Conclusion

In this paper, we focus on the challenge posed by highly sparse data matrices, which have become increasingly common in many areas, including biology and text mining. While much work has focused on addressing the challenges of high dimensional data, relatively little attention has been given to the challenges of sparsity in the data. We show, both theoretically and empirically, that not explicitly accounting for the sparsity in the data hurts one's prediction errors and one's ability to perform feature selection. Our proposed method is able to make productive use of highly sparse features by creating new aggregated features based on side information about the original features. In contrast to simpler tree-based aggregation strategies that are occasionally used as a pre-processing step in biological applications, our method adaptively learns the feature aggregation in a supervised manner. In doing so, our methodology not only overcomes the challenges of data sparsity but also produces features that may be of greater relevance to the particular prediction task of interest.

## Acknowledgments

The authors thank Andy Clark for calling our attention to the challenge of rare features. This work was supported by NSF CAREER grant, DMS-1653017.

## Appendices

### A Failure of OLS in the Presence of A Rare Feature

**Theorem 5.** *Consider the linear model (1) with  $\mathbf{X} \in \mathbb{R}^{n \times p}$  having full column rank. Further suppose that  $\mathbf{X}_j$  is a binary vector having  $k$  nonzeros. It follows that*

$$\mathbb{P} \left( \left| \hat{\beta}_j^{\text{OLS}}(n) - \beta_j^* \right| > \eta \right) \geq 2\Phi \left( -\eta k^{1/2} / \sigma \right) \quad \text{for any } \eta > 0, \quad (7)$$

where  $\Phi(\cdot)$  is the cumulative distribution function of a standard normal variable.

*Proof.* The distribution of the OLS estimator is  $\hat{\beta}_j^{\text{OLS}}(n) \sim N(\beta_j^*, \sigma^2[(\mathbf{X}^T \mathbf{X})^{-1}]_{jj})$ . By applying blockwise inversion (see, e.g., Bernstein 2009), with the  $j$ th row/column of  $\mathbf{X}^T \mathbf{X}$  in its own “block”, we get

$$\begin{aligned} [(\mathbf{X}^T \mathbf{X})^{-1}]_{jj} &= [\mathbf{X}_j^T \mathbf{X}_j - \mathbf{X}_{-j}^T \mathbf{X}_{-j} (\mathbf{X}_{-j}^T \mathbf{X}_{-j})^{-1} \mathbf{X}_{-j}^T \mathbf{X}_j]^{-1} \\ &= [\|\mathbf{X}_j\|^2 - \|(\mathbf{X}_{-j}^T \mathbf{X}_{-j})^{-1/2} \mathbf{X}_{-j}^T \mathbf{X}_j\|^2]^{-1} \\ &\geq \|\mathbf{X}_j\|^{-2} = k^{-1}. \end{aligned}$$

Thus,

$$\mathbb{P} \left( \left| \hat{\beta}_j^{\text{OLS}}(n) - \beta_j^* \right| > \eta \right) = 2\Phi \left( -\frac{\eta}{\sigma \sqrt{[(\mathbf{X}^T \mathbf{X})^{-1}]_{jj}}} \right) \geq 2\Phi \left( -\eta k^{1/2} / \sigma \right)$$

where  $\Phi(\cdot)$  is the distribution function of a standard normal variable.  $\square$

## B Proof of Theorem 2

In the setting of Theorem 2, we have  $\mathbf{X}\boldsymbol{\beta}^* = \widetilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}^*$ , where  $\widetilde{\mathbf{X}} = \mathbf{X}(\mathbf{I}_k \otimes \mathbf{1}_{p/k}) \in \mathbb{R}^{n \times k}$ . The two estimators, the oracle lasso on the aggregated data ( $\widetilde{\mathbf{X}}$ ) and the lasso on the original data ( $\mathbf{X}$ ), are defined below.

- Oracle lasso estimator  $\hat{\boldsymbol{\beta}}_\lambda^{oracle} = \check{\boldsymbol{\beta}}_\lambda^{oracle} \otimes \mathbf{1}_{p/k}$  where  $\check{\boldsymbol{\beta}}_\lambda^{oracle}$  is the unique solution to

$$\min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^k} \frac{1}{2n} \|\mathbf{y} - \widetilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}\|_2^2 + \lambda \|\tilde{\boldsymbol{\beta}}\|_1.$$

- Lasso estimator  $\hat{\boldsymbol{\beta}}_\lambda^{lasso}$  is defined in (2).

We begin by establishing that the interval  $\mathcal{I}$  is nonempty, which is ensured by the constraint  $k < p/(36 \log n)$ . In particular, the lower bound of the interval is below the upper bound if

$$12k \log(k^2 n) < p \log(2\tilde{c}p/k)$$

Now,  $\log(k^2 n) \leq 3 \log n$  and  $\log(2\tilde{c}p/k) \geq 1$  as long as  $2\tilde{c}p/k > e$ . Now,  $2\tilde{c} > 0.6$ , so if  $p/k \geq 5$  then it would suffice to show that

$$36k \log n < p.$$

And this constraint does imply that  $p/k \geq 5$ . The two parts of the theorem follow from the following two propositions.

**Proposition 2** (Support recovery of oracle lasso). *Suppose  $\min_{i=1, \dots, k-1} |\tilde{\beta}_i^*| > \sigma \sqrt{\frac{4k}{n} \log(k^2 n)}$ .*

*With  $\lambda = \sigma \sqrt{\frac{\log(k^2 n)}{kn}}$ , the oracle lasso recovers the correct signed support successfully:*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left( \mathbb{S}_\pm(\hat{\boldsymbol{\beta}}_\lambda^{oracle}) = \mathbb{S}_\pm(\boldsymbol{\beta}^*) \right) = 1.$$

*Proof.* The scaled matrix  $\sqrt{\frac{k}{n}}\widetilde{\mathbf{X}}$  is orthogonal since

$$\widetilde{\mathbf{X}}^T \widetilde{\mathbf{X}} = (\mathbf{I}_k \otimes \mathbf{1}_{p/k})^T \mathbf{X}^T \mathbf{X} (\mathbf{I}_k \otimes \mathbf{1}_{p/k}) = \frac{n}{p} \cdot \frac{p}{k} \mathbf{I}_k = \frac{n}{k} \mathbf{I}_k.$$

Orthogonality implies that

$$\check{\boldsymbol{\beta}}_\lambda^{oracle} = S \left( \left( \sqrt{\frac{k}{n}} \widetilde{\mathbf{X}}^T \right) \left( \sqrt{\frac{k}{n}} \mathbf{y} \right), \lambda k \right) = S \left( \frac{k}{n} \widetilde{\mathbf{X}}^T \mathbf{y}, \lambda k \right) \quad (8)$$

where  $\frac{k}{n} \widetilde{\mathbf{X}}^T \mathbf{y} = \frac{k}{n} \widetilde{\mathbf{X}}^T \widetilde{\mathbf{X}} \tilde{\boldsymbol{\beta}}^* + \frac{k}{n} \widetilde{\mathbf{X}}^T \boldsymbol{\varepsilon} = \tilde{\boldsymbol{\beta}}^* + \frac{k}{n} \widetilde{\mathbf{X}}^T \boldsymbol{\varepsilon} \sim N_k(\tilde{\boldsymbol{\beta}}^*, \frac{k\sigma^2}{n} \mathbf{I}_k)$ . By the Chernoff bound for normal variables, for any  $t > 0$ ,

$$\mathbb{P} \left( \left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} - \tilde{\beta}_j^* \right| > t \right) \leq 2 \exp \left( -\frac{t^2}{2k\sigma^2/n} \right) \quad \text{for } j = 1, \dots, k.$$

Choosing  $t = \sigma \sqrt{\frac{k \log(k^2 n)}{n}}$  and applying a union bound yields

$$\mathbb{P} \left( \left\| \frac{k}{n} \widetilde{\mathbf{X}}^T \mathbf{y} - \tilde{\boldsymbol{\beta}}^* \right\|_\infty > \sigma \sqrt{\frac{k \log(k^2 n)}{n}} \right) \leq 2k \exp \left( - \frac{\sigma^2 k \log(k^2 n)/n}{2k\sigma^2/n} \right) = \frac{2}{\sqrt{n}}.$$

Hence, with probability at least  $1 - \frac{2}{\sqrt{n}}$ , we have  $\left\| \frac{k}{n} \widetilde{\mathbf{X}}^T \mathbf{y} - \tilde{\boldsymbol{\beta}}^* \right\|_\infty \leq \sigma \sqrt{\frac{k \log(k^2 n)}{n}} = \lambda k$ , due to our choice of  $\lambda = \sigma \sqrt{\frac{\log(k^2 n)}{kn}}$ . Under  $\left\| \frac{k}{n} \widetilde{\mathbf{X}}^T \mathbf{y} - \tilde{\boldsymbol{\beta}}^* \right\|_\infty \leq \lambda k$ , the following results hold.

- By  $\tilde{\beta}_k^* = 0$  and

$$\left| \frac{k}{n} (\widetilde{\mathbf{X}}_k)^T \mathbf{y} \right| = \left| \frac{k}{n} (\widetilde{\mathbf{X}}_k)^T \mathbf{y} - \tilde{\beta}_k^* \right| \leq \left\| \frac{k}{n} \widetilde{\mathbf{X}}^T \mathbf{y} - \tilde{\boldsymbol{\beta}}^* \right\|_\infty \leq \lambda k,$$

we have  $\check{\beta}_{\lambda,k}^{oracle} = S \left( \frac{k}{n} (\widetilde{\mathbf{X}}_k)^T \mathbf{y}, \lambda k \right) = 0$  and  $\check{\beta}_{\lambda,\ell}^{oracle} = \check{\beta}_{\lambda,k}^{oracle} = \tilde{\beta}_k^* = \beta_\ell^*$  for all  $\ell > \frac{k-1}{k}n$ .

- For  $j = 1, \dots, k-1$ , since  $\left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} - \tilde{\beta}_j^* \right| \leq \lambda k$  and  $\left| \tilde{\beta}_j^* \right| \geq \min_{i=1, \dots, k-1} \left| \tilde{\beta}_i^* \right| > 2\lambda k$ , we must have  $\frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y}$  and  $\tilde{\beta}_j^*$  share the same sign. Moreover, we either have

$$\left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} \right| \geq \left| \tilde{\beta}_j^* \right| > \lambda k$$

or  $\left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} \right| < \left| \tilde{\beta}_j^* \right|$  in which case  $\left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} - \tilde{\beta}_j^* \right| = \left| \left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} \right| - \left| \tilde{\beta}_j^* \right| \right| = \left| \tilde{\beta}_j^* \right| - \left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} \right| \leq \lambda k$  and therefore

$$\left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} \right| \geq \left| \tilde{\beta}_j^* \right| - \lambda k \geq 2\lambda k - \lambda k = \lambda k.$$

Thus,  $\left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} \right| > \lambda k$  for  $j = 1, \dots, k-1$ . By definition of  $\hat{\boldsymbol{\beta}}_\lambda^{oracle}$  and (8), for  $\frac{j-1}{k}n < \ell \leq \frac{j}{k}n$ ,

$$\hat{\beta}_{\lambda,\ell}^{oracle} = \check{\beta}_{\lambda,j}^{oracle} = S \left( \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y}, \lambda k \right) = \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} \left( 1 - \frac{\lambda k}{\left| \frac{k}{n} (\widetilde{\mathbf{X}}_j)^T \mathbf{y} \right|} \right)$$

which is of the same sign as  $\tilde{\beta}_j^*$  (and the same sign as  $\beta_\ell^*$ ).

In the above two bullet points, we have shown  $\mathbb{S}_\pm(\hat{\boldsymbol{\beta}}_\lambda^{oracle}) = \mathbb{S}_\pm(\boldsymbol{\beta}^*)$  holds with probability at least  $1 - \frac{2}{\sqrt{n}}$ . Hence,

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left( \mathbb{S}_\pm(\hat{\boldsymbol{\beta}}_\lambda^{oracle}) = \mathbb{S}_\pm(\boldsymbol{\beta}^*) \right) \geq \lim_{n \rightarrow \infty} 1 - \frac{2}{\sqrt{n}} = 1.$$

Since  $\limsup_{n \rightarrow \infty} \mathbb{P} \left( \mathbb{S}_\pm(\hat{\boldsymbol{\beta}}_\lambda^{oracle}) = \mathbb{S}_\pm(\boldsymbol{\beta}^*) \right) = 1$ , the limit for  $\mathbb{P} \left( \mathbb{S}_\pm(\hat{\boldsymbol{\beta}}_\lambda^{oracle}) = \mathbb{S}_\pm(\boldsymbol{\beta}^*) \right)$  is 1.  $\square$

**Lemma 2.** Suppose  $\boldsymbol{\varepsilon} \sim N_n(\mathbf{0}, \sigma^2 \mathbf{I}_n)$  and  $\tilde{c} = \frac{1}{3} e^{(\pi/2+2)^{-1}} \sqrt{\frac{1}{4} + \frac{1}{\pi}}$ . Then

$$\mathbb{P} \left( \max_{\ell=1,\dots,p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| \leq 2\sigma \sqrt{\frac{n}{3p} \log(2\tilde{c}p)} \right) \leq \left( 1 - \frac{1}{p} \right)^p.$$

*Proof.* Let  $Z$  be a standard Gaussian variable. Theorem 2.1 of Côté et al. (2012) provides a lower bound for the Gaussian Q function (i.e.,  $\mathbb{P}(Z > z)$ ). Choosing  $\kappa = \frac{3}{2}$  in their Theorem 2.1 yields

$$\mathbb{P}(Z > z) \geq \underbrace{\left( \frac{1}{3} e^{(\pi/2+2)^{-1}} \sqrt{\frac{1}{4} + \frac{1}{\pi}} \right)}_{\tilde{c}} e^{-\frac{3z^2}{4}}$$

where  $\tilde{c} = \frac{1}{3} e^{(\pi/2+2)^{-1}} \sqrt{\frac{1}{4} + \frac{1}{\pi}}$  is independent of  $z$ . Given that  $\mathbf{X}_\ell$  has  $n/p$  one's and  $\mathbf{X}_1 \mathbf{1}_p = \mathbf{1}_n$ , we have  $\mathbf{X}_\ell^T \boldsymbol{\varepsilon} \stackrel{iid}{\sim} N(0, \frac{n}{p} \sigma^2)$  for  $\ell = 1, \dots, p$ . By expressing  $\mathbf{X}_1^T \boldsymbol{\varepsilon} = \sqrt{\frac{n}{p}} \sigma Z$ , we have for any  $\eta > 0$ ,

$$\mathbb{P}(\mathbf{X}_1^T \boldsymbol{\varepsilon} > \eta) \geq \tilde{c} e^{-\frac{3\eta^2 p}{4\sigma^2 n}} \quad \Rightarrow \quad \mathbb{P}(|\mathbf{X}_1^T \boldsymbol{\varepsilon}| > \eta) \geq 2\tilde{c} e^{-\frac{3\eta^2 p}{4\sigma^2 n}}.$$

Moreover,

$$\mathbb{P} \left( \max_{\ell=1,\dots,p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| \leq \eta \right) = (\mathbb{P}(|\mathbf{X}_1^T \boldsymbol{\varepsilon}| \leq \eta))^p = (1 - \mathbb{P}(|\mathbf{X}_1^T \boldsymbol{\varepsilon}| > \eta))^p \leq \left( 1 - 2\tilde{c} e^{-\frac{3\eta^2 p}{4\sigma^2 n}} \right)^p.$$

Plugging in  $\eta = 2\sigma \sqrt{\frac{n}{3p} \log(2\tilde{c}p)}$  in the above inequality yields

$$\mathbb{P} \left( \max_{\ell=1,\dots,p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| \leq 2\sigma \sqrt{\frac{n}{3p} \log(2\tilde{c}p)} \right) \leq \left( 1 - \frac{1}{p} \right)^p.$$

□

**Proposition 3** (Failure of support recovery of lasso). Suppose  $\min_{i=1,\dots,k-1} |\tilde{\beta}_i^*| \leq \sigma \sqrt{\frac{p}{3n} \log(2\tilde{c}p/k)}$  where  $\tilde{c} = \frac{1}{3} e^{(\pi/2+2)^{-1}} \sqrt{\frac{1}{4} + \frac{1}{\pi}}$ . The lasso fails to get high-probability signed support recovery:

$$\limsup_{p \rightarrow \infty} \sup_{\lambda \geq 0} \mathbb{P} \left( \mathbb{S}_\pm(\hat{\beta}_\lambda^{lasso}) = \mathbb{S}_\pm(\boldsymbol{\beta}^*) \right) \leq \frac{1}{e}.$$

*Proof.* The lasso solution can be simplified to  $\hat{\beta}_\lambda^{lasso} = S(\frac{p}{n} \mathbf{X}^T \mathbf{y}, \lambda p)$ . Since  $\beta_\ell^* \neq 0$  for  $\ell \leq \frac{k-1}{k}p$  and  $\beta_\ell^* = 0$  for  $\ell > \frac{k-1}{k}p$ , the following is a necessary condition for  $\hat{\beta}_\lambda^{lasso}$  to recover the correct signed support:

$$\exists \lambda \text{ s.t. } |\mathbf{X}_\ell^T \mathbf{y}| > \lambda p \text{ for } \ell \leq \frac{k-1}{k}p \text{ and } |\mathbf{X}_\ell^T \mathbf{y}| \leq \lambda p \text{ for } \ell > \frac{k-1}{k}p$$

↔

$$\min_{\ell \leq \frac{k-1}{k}p} |\mathbf{X}_\ell^T \mathbf{y}| > \max_{\ell > \frac{k-1}{k}p} |\mathbf{X}_\ell^T \mathbf{y}|.$$

Furthermore, we have

$$\mathbf{X}_\ell^T \mathbf{y} = \mathbf{X}_\ell^T \left( \sum_{j=1}^p \mathbf{X}_j \beta_j^* + \boldsymbol{\varepsilon} \right) = \frac{n}{p} \beta_\ell^* + \mathbf{X}_\ell^T \boldsymbol{\varepsilon}.$$

Define  $\bar{i} := \arg \min_{i=1, \dots, k-1} |\tilde{\beta}_i^*|$  and  $\mathcal{A} := \left\{ \max_{\ell > \frac{k-1}{k}p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| \leq 2\sigma \sqrt{\frac{n}{3p} \log(2\tilde{c}p/k)} \right\}$ . Then

$$\begin{aligned} & \mathbb{P} \left( \mathbb{S}_\pm(\hat{\boldsymbol{\beta}}_\lambda^{lasso}) = \mathbb{S}_\pm(\boldsymbol{\beta}^*) \right) \\ & \leq \mathbb{P} \left( \min_{\ell \leq \frac{k-1}{k}p} |\mathbf{X}_\ell^T \mathbf{y}| > \max_{\ell > \frac{k-1}{k}p} |\mathbf{X}_\ell^T \mathbf{y}| \right) \\ & \leq \mathbb{P} \left( \min_{\frac{\bar{i}-1}{k}p < \ell \leq \frac{\bar{i}}{k}p} |\mathbf{X}_\ell^T \mathbf{y}| > \max_{\ell > \frac{k-1}{k}p} |\mathbf{X}_\ell^T \mathbf{y}| \right) \\ & \leq \mathbb{P} \left( \frac{n}{p} |\tilde{\beta}_{\bar{i}}^*| + \min_{\frac{\bar{i}-1}{k}p < \ell \leq \frac{\bar{i}}{k}p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| > \max_{\ell > \frac{k-1}{k}p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| \right) \\ & = \mathbb{P} \left( \frac{n}{p} |\tilde{\beta}_{\bar{i}}^*| + \min_{\frac{\bar{i}-1}{k}p < \ell \leq \frac{\bar{i}}{k}p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| > \max_{\ell > \frac{k-1}{k}p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| \text{ and } \mathcal{A}^c \right) \\ & \quad + \mathbb{P} \left( \frac{n}{p} |\tilde{\beta}_{\bar{i}}^*| + \min_{\frac{\bar{i}-1}{k}p < \ell \leq \frac{\bar{i}}{k}p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| > \max_{\ell > \frac{k-1}{k}p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| \middle| \mathcal{A} \right) \cdot \mathbb{P}(\mathcal{A}) \\ & \leq \mathbb{P} \left( \frac{n}{p} |\tilde{\beta}_{\bar{i}}^*| + \min_{\frac{\bar{i}-1}{k}p < \ell \leq \frac{\bar{i}}{k}p} |\mathbf{X}_\ell^T \boldsymbol{\varepsilon}| > 2\sigma \sqrt{\frac{n}{3p} \log(2\tilde{c}p/k)} \right) + \mathbb{P}(\mathcal{A}) \\ & = \left[ \mathbb{P} \left( |\mathbf{X}_1^T \boldsymbol{\varepsilon}| > 2\sigma \sqrt{\frac{n}{3p} \log(2\tilde{c}p/k)} - \frac{n}{p} |\tilde{\beta}_{\bar{i}}^*| \right) \right]^{\frac{p}{k}} + \mathbb{P}(\mathcal{A}) \quad (\text{by } \mathbf{X}_\ell^T \boldsymbol{\varepsilon} \text{ being } i.i.d.) \\ & \leq \left[ \mathbb{P} \left( |\mathbf{X}_1^T \boldsymbol{\varepsilon}| > \sigma \sqrt{\frac{n}{3p} \log(2\tilde{c}p/k)} \right) \right]^{\frac{p}{k}} + \mathbb{P}(\mathcal{A}) \quad \left( \text{by } |\tilde{\beta}_{\bar{i}}^*| \leq \sigma \sqrt{\frac{p}{3n} \log(2\tilde{c}p/k)} \right) \\ & \leq \left[ 2 \exp \left( -\frac{p}{2\sigma^2 n} \cdot \frac{\sigma^2 n}{3p} \log(2\tilde{c}p/k) \right) \right]^{\frac{p}{k}} + \left( 1 - \frac{k}{p} \right)^{\frac{p}{k}} \quad (\text{by Chernoff ineq and Lemma 2}) \\ & \leq 2^{\frac{p}{k}} \exp \left( -\frac{p}{6k} \log(2\tilde{c}p/k) \right) + \left( 1 - \frac{k}{p} \right)^{\frac{p}{k}} \\ & = 2^{\frac{p}{k}} \left( \frac{2\tilde{c}p}{k} \right)^{-\frac{p}{6k}} + \left( 1 - \frac{k}{p} \right)^{\frac{p}{k}} \\ & = \left( \frac{\tilde{c}p}{32k} \right)^{-\frac{p}{6k}} + \left( 1 - \frac{k}{p} \right)^{\frac{p}{k}} \end{aligned}$$

which holds for all  $\lambda \geq 0$ . In particular,

$$\sup_{\lambda \geq 0} \mathbb{P} \left( \mathbb{S}_\pm(\hat{\boldsymbol{\beta}}_\lambda^{lasso}) = \mathbb{S}_\pm(\boldsymbol{\beta}^*) \right) \leq \left( \frac{\tilde{c}p}{32k} \right)^{-\frac{p}{6k}} + \left( 1 - \frac{k}{p} \right)^{\frac{p}{k}}.$$

Taking  $\limsup$  on both side yields

$$\begin{aligned} \limsup_{p \rightarrow \infty} \sup_{\lambda \geq 0} \mathbb{P} \left( \mathbb{S}_{\pm}(\hat{\beta}_{\lambda}^{lasso}) = \mathbb{S}_{\pm}(\beta^*) \right) &\leq \lim_{p \rightarrow \infty} \left( \frac{\tilde{c}p}{32k} \right)^{-\frac{p}{6k}} + \lim_{p \rightarrow \infty} \left( 1 - \frac{k}{p} \right)^{\frac{p}{k}} \\ &= 0 + \frac{1}{e} = \frac{1}{e}. \end{aligned}$$

□

## C Proof of Theorem 3

Let  $\mathbf{A} \in \{0, 1\}^{p \times |\mathcal{C}|}$  denote the aggregation matrix corresponding to the partition  $\mathcal{C}$ . Aggregated least squares uses the design matrix  $\widetilde{\mathbf{X}} = \mathbf{X}\mathbf{A}$ , where each column is the sum of features in a group of  $\mathcal{C}$ . Left-multiplication by  $\mathbf{A}$  maps this back to  $p$ -dimensional space:

$$\hat{\beta}^{\mathcal{C}} = \mathbf{A}\widetilde{\mathbf{X}}^+\mathbf{y}.$$

Its fitted values are

$$\mathbf{X}\hat{\beta}^{\mathcal{C}} = \widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+\mathbf{y} = \widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+(\mathbf{X}\beta^* + \varepsilon).$$

Now,  $\widetilde{\mathbf{X}}$  is full-rank and

$$\mathbf{X}\hat{\beta}^{\mathcal{C}} \sim N_p(\widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+\mathbf{X}\beta^*, \sigma^2\widetilde{\mathbf{X}}[\widetilde{\mathbf{X}}^T\widetilde{\mathbf{X}}]^{-1}\widetilde{\mathbf{X}}^T).$$

Now,

$$\begin{aligned} \mathbb{E}[\mathbf{X}\hat{\beta}^{\mathcal{C}}] &= \widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+\mathbf{X}\beta^* \\ &= \widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+\mathbf{X}\mathbf{A}\mathbf{A}^+\beta^* + \widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+\mathbf{X}(\mathbf{I}_p - \mathbf{A}\mathbf{A}^+)\beta^* \\ &= \mathbf{X}\mathbf{A}\mathbf{A}^+\beta^* + \widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+\mathbf{X}(\mathbf{I}_p - \mathbf{A}\mathbf{A}^+)\beta^* \end{aligned}$$

since  $\mathbf{X}\mathbf{A}\mathbf{A}^+\beta^*$  is in the column space of  $\widetilde{\mathbf{X}}$ , and

$$\begin{aligned} \mathbb{E}[\mathbf{X}\hat{\beta}^{\mathcal{C}}] - \mathbf{X}\beta^* &= \mathbf{X}(\mathbf{A}\mathbf{A}^+ - \mathbf{I}_p)\beta^* + \widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+\mathbf{X}(\mathbf{I}_p - \mathbf{A}\mathbf{A}^+)\beta^* \\ &= (\widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+ - \mathbf{I}_n)\mathbf{X}(\mathbf{I}_p - \mathbf{A}\mathbf{A}^+)\beta^* \end{aligned}$$

$$\begin{aligned} \mathbb{E}\|\mathbf{X}\hat{\beta}^{\mathcal{C}} - \mathbf{X}\beta^*\|^2 &= \|(\widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+ - \mathbf{I}_n)\mathbf{X}(\mathbf{I}_p - \mathbf{A}\mathbf{A}^+)\beta^*\|^2 + \sigma^2\text{tr}(\widetilde{\mathbf{X}}[\widetilde{\mathbf{X}}^T\widetilde{\mathbf{X}}]^{-1}\widetilde{\mathbf{X}}^T) \\ &= \|(\widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+ - \mathbf{I}_n)\mathbf{X}(\mathbf{I}_p - \mathbf{A}\mathbf{A}^+)\beta^*\|^2 + \sigma^2|\mathcal{C}| \\ &\leq \|\mathbf{X}\|_{op}^2\|(\mathbf{I}_p - \mathbf{A}\mathbf{A}^+)\beta^*\|^2 + \sigma^2|\mathcal{C}| \end{aligned}$$

since  $\|(\widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^+ - \mathbf{I}_n)\|_{op} = 1$  and the trace of a projection matrix is the rank of the target space. Finally, observe that

$$\|(\mathbf{I}_p - \mathbf{A}\mathbf{A}^+)\beta^*\|^2 = \sum_{\ell=1}^{|\mathcal{C}|} \sum_{j \in \mathcal{C}_{\ell}} \left( \beta_j^* - |\mathcal{C}_{\ell}|^{-1} \sum_{j' \in \mathcal{C}_{\ell}} \beta_{j'}^* \right)^2.$$

## D Consensus ADMM for Solving Problem (5)

The ADMM algorithm is given in Algorithm 1. Let  $\mathbf{X} = \text{SVD}_{\text{compact}}(\tilde{\mathbf{U}}, \tilde{\mathbf{D}}, \tilde{\mathbf{V}})$  be the *compact* singular value decomposition of  $\mathbf{X}$ , where  $\tilde{\mathbf{D}} \in \mathbb{R}^{\min(n,p) \times \min(n,p)}$  is a diagonal matrix with non-zero singular values on the diagonal, and  $\tilde{\mathbf{U}} \in \mathbb{R}^{n \times \min(n,p)}$  and  $\tilde{\mathbf{V}} \in \mathbb{R}^{p \times \min(n,p)}$  contain the left and right singular vectors in columns corresponding to non-zero singular values, respectively. Similarly, we have  $(\mathbf{I}_p : -\mathbf{A}) = \text{SVD}_{\text{compact}}(\cdot, \cdot, \tilde{\mathbf{Q}})$  where  $\tilde{\mathbf{Q}} \in \mathbb{R}^{(p+|\mathcal{T}|) \times p}$  contains  $p$  right singular vectors corresponding to non-zero singular values.

---

**Algorithm 1** *Consensus ADMM for Solving Problem (5)*

---

**Input:**  $\mathbf{y}, \mathbf{X}, \mathbf{A}, n, p, |\mathcal{T}|, \lambda, \alpha, \rho, \epsilon^{abs}, \epsilon^{rel}, \text{maxite}$ .

```

1:  $\mathbf{X} = \text{SVD}_{\text{compact}}(\cdot, \tilde{\mathbf{D}}, \tilde{\mathbf{V}})$ 
2:  $(\mathbf{I}_p : -\mathbf{A}) = \text{SVD}_{\text{compact}}(\cdot, \cdot, \tilde{\mathbf{Q}})$ 
3:  $\boldsymbol{\beta}^0 \leftarrow \boldsymbol{\beta}^{(i)0} \leftarrow \mathbf{v}^{(i)0} \leftarrow \mathbf{0} \in \mathbb{R}^p \quad \forall i = 1, 2, 3$ 
4:  $\boldsymbol{\gamma}^0 \leftarrow \boldsymbol{\gamma}^{(j)0} \leftarrow \mathbf{u}^{(j)0} \leftarrow \mathbf{0} \in \mathbb{R}^{|\mathcal{T}|} \quad \forall j = 1, 2$ 
5: continue  $\leftarrow$  true
6:  $k \leftarrow 0$ 
7: while  $k < \text{maxite}$  and continue do
8:    $k \leftarrow k + 1$ 
9:    $\boldsymbol{\beta}^{(1)k} \leftarrow \left[ \tilde{\mathbf{V}} \text{diag} \left( \frac{1}{[\tilde{\mathbf{D}}^T \tilde{\mathbf{D}}]_{ii} + n\rho} \right) \tilde{\mathbf{V}}^T + \frac{1}{n\rho} (\mathbf{I}_p - \tilde{\mathbf{V}} \tilde{\mathbf{V}}^T) \right] (\mathbf{X}^T \mathbf{y} + n\rho \boldsymbol{\beta}^{k-1} - n\mathbf{v}^{(1)k-1})$ 
10:   $\beta_j^{(2)k} \leftarrow S \left( \beta_j^{k-1} - \frac{1}{\rho} v_j^{(2)k-1}, \frac{\lambda(1-\alpha)\bar{w}_j}{\rho} \right) \quad \forall j = 1, \dots, p$ 
11:   $\gamma_\ell^{(1)k} \leftarrow \begin{cases} S \left( \gamma_\ell^{k-1} - \frac{1}{\rho} u_\ell^{(1)k-1}, \frac{\lambda\alpha w_\ell}{\rho} \right) & \text{if } w_\ell > 0 \\ \gamma_\ell^{k-1} - \frac{1}{\rho} u_\ell^{(1)k-1} & \text{if } w_\ell = 0 \end{cases}$ 
12:   $\begin{pmatrix} \boldsymbol{\beta}^{(3)k} \\ \boldsymbol{\gamma}^{(2)k} \end{pmatrix} \leftarrow (\mathbf{I}_{p+|\mathcal{T}|} - \tilde{\mathbf{Q}} \tilde{\mathbf{Q}}^T) \begin{pmatrix} \boldsymbol{\beta}^{k-1} \\ \boldsymbol{\gamma}^{k-1} \end{pmatrix} - \frac{1}{\rho} \begin{pmatrix} \mathbf{v}^{(3)k-1} \\ \mathbf{u}^{(2)k-1} \end{pmatrix}$ 
13:   $\boldsymbol{\beta}^k \leftarrow (\boldsymbol{\beta}^{(1)k} + \boldsymbol{\beta}^{(2)k} + \boldsymbol{\beta}^{(3)k})/3$ 
14:   $\boldsymbol{\gamma}^k \leftarrow (\boldsymbol{\gamma}^{(1)k} + \boldsymbol{\gamma}^{(2)k})/2$ 
15:   $\mathbf{v}^{(i)k} \leftarrow \mathbf{v}^{(i)k-1} + \rho(\boldsymbol{\beta}^{(i)k} - \boldsymbol{\beta}^k) \quad \forall i = 1, 2, 3$ 
16:   $\mathbf{u}^{(j)k} \leftarrow \mathbf{u}^{(j)k-1} + \rho(\boldsymbol{\gamma}^{(j)k} - \boldsymbol{\gamma}^k) \quad \forall j = 1, 2$ 
17:  if  $\sqrt{\sum_{i=1}^3 \|\boldsymbol{\beta}^{(i)k} - \boldsymbol{\beta}^k\|_2^2 + \sum_{j=1}^2 \|\boldsymbol{\gamma}^{(j)k} - \boldsymbol{\gamma}^k\|_2^2} \leq \epsilon^{abs} \sqrt{3p + 2|\mathcal{T}|} + \epsilon^{rel} \max \left\{ \sqrt{\sum_{i=1}^3 \|\boldsymbol{\beta}^{(i)k}\|_2^2 + \sum_{j=1}^2 \|\boldsymbol{\gamma}^{(j)k}\|_2^2}, \sqrt{3 \|\boldsymbol{\beta}^k\|_2^2 + 2 \|\boldsymbol{\gamma}^k\|_2^2} \right\}$ 
   and  $\rho \sqrt{3 \|\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k-1}\|_2^2 + 2 \|\boldsymbol{\gamma}^k - \boldsymbol{\gamma}^{k-1}\|_2^2} \leq \epsilon^{abs} \sqrt{3p + 2|\mathcal{T}|} + \epsilon^{rel} \sqrt{\sum_{i=1}^3 \|\mathbf{v}^{(i)k}\|_2^2 + \sum_{j=1}^2 \|\mathbf{u}^{(j)k}\|_2^2}$  then
18:    continue  $\leftarrow$  false
19:  end if
20: end while

```

**Output:**  $\boldsymbol{\beta}^k, \boldsymbol{\gamma}^k$

---

## D.1 Derivation of Algorithm 1

The ADMM updates involve minimizing the augmented Lagrangian of the global consensus problem (6),

$$\begin{aligned}
& L_\rho(\boldsymbol{\beta}^{(1)}, \boldsymbol{\beta}^{(2)}, \boldsymbol{\beta}^{(3)}, \boldsymbol{\gamma}^{(1)}, \boldsymbol{\gamma}^{(2)}, \boldsymbol{\beta}, \boldsymbol{\gamma}; \mathbf{v}^{(1)}, \mathbf{v}^{(2)}, \mathbf{v}^{(3)}, \mathbf{u}^{(1)}, \mathbf{u}^{(2)}) \\
&= \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(1)}\|_2^2 + \lambda \left( \alpha \sum_{\ell=1}^{|\mathcal{T}|} w_\ell \left| \gamma_\ell^{(1)} \right| + (1-\alpha) \sum_{j=1}^p \tilde{w}_j \left| \beta_j^{(2)} \right| \right) + 1_\infty \{ \boldsymbol{\beta}^{(3)} = A\boldsymbol{\gamma}^{(2)} \} \\
&\quad + \sum_{i=1}^3 \left( \mathbf{v}^{(i)T} (\boldsymbol{\beta}^{(i)} - \boldsymbol{\beta}) + \frac{\rho}{2} \|\boldsymbol{\beta}^{(i)} - \boldsymbol{\beta}\|_2^2 \right) + \sum_{j=1}^2 \left( \mathbf{u}^{(j)T} (\boldsymbol{\gamma}^{(j)} - \boldsymbol{\gamma}) + \frac{\rho}{2} \|\boldsymbol{\gamma}^{(j)} - \boldsymbol{\gamma}\|_2^2 \right).
\end{aligned}$$

1. Update  $\boldsymbol{\beta}^{(1)}$ .

$$\boldsymbol{\beta}^{(1)k+1} := \arg \min_{\boldsymbol{\beta}^{(1)} \in \mathbb{R}^p} \left\{ \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(1)}\|_2^2 + \langle \mathbf{v}^{(1)k}, (\boldsymbol{\beta}^{(1)} - \boldsymbol{\beta}^k) \rangle + \frac{\rho}{2} \|\boldsymbol{\beta}^{(1)} - \boldsymbol{\beta}^k\|_2^2 \right\}.$$

Let  $\mathbf{X} = \text{SVD}(\mathbf{U}, \mathbf{D}, \mathbf{V})$  be the singular value decomposition of  $\mathbf{X}$ , where  $\mathbf{U} \in \mathbb{R}^{n \times n}$  contains left singular vectors in columns,  $\mathbf{V} \in \mathbb{R}^{p \times p}$  contains right singular vectors in columns, and  $\mathbf{D} \in \mathbb{R}^{n \times p}$  is a rectangular diagonal matrix with decreasing singular values on the diagonal. First order condition to the above problem gives us:

$$\begin{aligned}
& (\mathbf{X}^T \mathbf{X} + n\rho \mathbf{I}_p) \boldsymbol{\beta}^{(1)k+1} = \mathbf{X}^T \mathbf{y} + n\rho \boldsymbol{\beta}^k - n\mathbf{v}^{(1)k} \\
& \Rightarrow \mathbf{V}(\mathbf{D}^T \mathbf{D} + n\rho \mathbf{I}_p) \mathbf{V}^T \boldsymbol{\beta}^{(1)k+1} = \mathbf{X}^T \mathbf{y} + n\rho \boldsymbol{\beta}^k - n\mathbf{v}^{(1)k} \\
& \Rightarrow \boldsymbol{\beta}^{(1)k+1} = \mathbf{V} \text{diag} \left( ([\mathbf{D}^T \mathbf{D}]_{ii} + n\rho)^{-1} \right) \mathbf{V}^T (\mathbf{X}^T \mathbf{y} + n\rho \boldsymbol{\beta}^k - n\mathbf{v}^{(1)k}).
\end{aligned}$$

When  $n \geq p$ , we have

$$\boldsymbol{\beta}^{(1)k+1} = \tilde{\mathbf{V}} \text{diag} \left( ([\tilde{\mathbf{D}}^T \tilde{\mathbf{D}}]_{ii} + n\rho)^{-1} \right) \tilde{\mathbf{V}}^T (\mathbf{X}^T \mathbf{y} + n\rho \boldsymbol{\beta}^k - n\mathbf{v}^{(1)k}). \quad (9)$$

When  $n < p$ , the SVD can be expressed in a compact form:  $\mathbf{D} = (\tilde{\mathbf{D}} : \mathbf{0})$  and  $\mathbf{V} = (\tilde{\mathbf{V}} : \tilde{\mathbf{V}}_\perp)$  where  $\tilde{\mathbf{D}} \in \mathbb{R}^{n \times n}$  and  $\tilde{\mathbf{V}} \in \mathbb{R}^{p \times n}$  are from the compact SVD of  $\mathbf{X}$ , and  $\tilde{\mathbf{V}}_\perp \in \mathbb{R}^{p \times (p-n)}$ . Thus,

$$\begin{aligned}
\mathbf{V} \text{diag} \left( ([\mathbf{D}^T \mathbf{D}]_{ii} + n\rho)^{-1} \right) \mathbf{V}^T &= \left( \tilde{\mathbf{V}} : \tilde{\mathbf{V}}_\perp \right) \text{diag} \left( ([\mathbf{D}^T \mathbf{D}]_{ii} + n\rho)^{-1} \right) \begin{pmatrix} \tilde{\mathbf{V}}^T \\ \tilde{\mathbf{V}}_\perp^T \end{pmatrix} \\
&= \tilde{\mathbf{V}} \text{diag} \left( ([\tilde{\mathbf{D}}^T \tilde{\mathbf{D}}]_{ii} + n\rho)^{-1} \right) \tilde{\mathbf{V}}^T + \tilde{\mathbf{V}}_\perp \tilde{\mathbf{V}}_\perp^T / (n\rho) \\
&= \tilde{\mathbf{V}} \text{diag} \left( ([\tilde{\mathbf{D}}^T \tilde{\mathbf{D}}]_{ii} + n\rho)^{-1} \right) \tilde{\mathbf{V}}^T + (\mathbf{I}_p - \tilde{\mathbf{V}} \tilde{\mathbf{V}}^T) / (n\rho).
\end{aligned}$$

So when  $n < p$ ,

$$\boldsymbol{\beta}^{(1)k+1} = \left[ \tilde{\mathbf{V}} \text{diag} \left( ([\tilde{\mathbf{D}}^T \tilde{\mathbf{D}}]_{ii} + n\rho)^{-1} \right) \tilde{\mathbf{V}}^T + (\mathbf{I}_p - \tilde{\mathbf{V}} \tilde{\mathbf{V}}^T) / (n\rho) \right] (\mathbf{X}^T \mathbf{y} + n\rho \boldsymbol{\beta}^k - n\mathbf{v}^{(1)k}). \quad (10)$$

Since  $\tilde{\mathbf{V}} = \mathbf{V}$  when  $n \geq p$  and  $\mathbf{V}\mathbf{V}^T = \mathbf{I}_p$ , we have (10) boil to (9) in that case.

2. Update  $\beta^{(2)}$ .

$$\beta^{(2)k+1} := \arg \min_{\beta^{(2)} \in \mathbb{R}^p} \left\{ \frac{\rho}{2} \left\| \beta^{(2)} - \left( \beta^k - \frac{1}{\rho} \mathbf{v}^{(2)k} \right) \right\|_2^2 + \lambda(1-\alpha) \sum_{j=1}^p \tilde{w}_j \left| \beta_j^{(2)} \right| \right\}.$$

The solution is simply elementwise soft-thresholding:

$$\beta_j^{(2)k+1} = S \left( \beta_j^k - \frac{1}{\rho} v_j^{(2)}, \frac{\lambda(1-\alpha)\tilde{w}_j}{\rho} \right) \quad \forall j = 1, \dots, p.$$

3. Update  $\gamma^{(1)}$ .

$$\gamma^{(1)k+1} := \arg \min_{\gamma^{(1)} \in \mathbb{R}^{|\mathcal{T}|}} \left\{ \frac{\rho}{2} \left\| \gamma^{(1)} - \left( \gamma^k - \frac{1}{\rho} \mathbf{u}^{(1)k} \right) \right\|_2^2 + \lambda \alpha \sum_{\ell=1}^{|\mathcal{T}|} w_\ell \left| \gamma_\ell^{(1)} \right| \right\}.$$

Since sometimes we will choose to  $w_\ell = 0$  for the root, we break the solution into cases:

$$\gamma_\ell^{(1)k+1} = \begin{cases} S \left( \gamma_\ell^k - \frac{1}{\rho} u_\ell^{(1)k}, \frac{\lambda \alpha w_\ell}{\rho} \right) & \text{if } w_\ell > 0 \\ \gamma_\ell^k - \frac{1}{\rho} u_\ell^{(1)k} & \text{if } w_\ell = 0. \end{cases}$$

4. Joint update of  $\beta^{(3)}$  and  $\gamma^{(2)}$ .

$$\begin{aligned} \begin{pmatrix} \beta^{(3)k+1} \\ \gamma^{(2)k+1} \end{pmatrix} &:= \arg \min_{\beta^{(3)} \in \mathbb{R}^p, \gamma^{(2)} \in \mathbb{R}^{|\mathcal{T}|}} \left\{ \left\| \beta^{(3)} - \left( \beta^k - \frac{1}{\rho} \mathbf{v}^{(3)k} \right) \right\|_2^2 + \left\| \gamma^{(2)} - \left( \gamma^k - \frac{1}{\rho} \mathbf{u}^{(2)k} \right) \right\|_2^2 \right\} \\ \text{s.t. } (\mathbf{I}_p : -\mathbf{A}) \begin{pmatrix} \beta^{(3)} \\ \gamma^{(2)} \end{pmatrix} &= 0. \end{aligned}$$

The solution is the projection of  $\begin{pmatrix} \beta^k \\ \gamma^k \end{pmatrix} - \frac{1}{\rho} \begin{pmatrix} \mathbf{v}^{(3)k} \\ \mathbf{u}^{(2)k} \end{pmatrix}$  onto the null space of  $(\mathbf{I}_p : -\mathbf{A})$ .

Let  $(\mathbf{I}_p : -\mathbf{A}) = \text{SVD}(\cdot, \mathbf{Q})$  where  $\mathbf{Q} = (\tilde{\mathbf{Q}} : \tilde{\mathbf{Q}}_\perp) \in \mathbb{R}^{(p+|\mathcal{T}|):(p+|\mathcal{T}|)}$  contains all the right singular vectors in columns. So  $\mathbf{I}_{p+|\mathcal{T}|} = \mathbf{Q} \mathbf{Q}^T = \tilde{\mathbf{Q}} \tilde{\mathbf{Q}}^T + \tilde{\mathbf{Q}}_\perp \tilde{\mathbf{Q}}_\perp^T$ . Since  $\tilde{\mathbf{Q}}$  corresponds to non-zero singular values of  $(\mathbf{I}_p : -\mathbf{A})$  by construction, we have  $\tilde{\mathbf{Q}}_\perp$  corresponds to the zero singular values, making itself an orthonormal basis for the null space of  $(\mathbf{I}_p : -\mathbf{A})$ . Thus,

$$\begin{aligned} \begin{pmatrix} \beta^{(3)k+1} \\ \gamma^{(2)k+1} \end{pmatrix} &= \tilde{\mathbf{Q}}_\perp (\tilde{\mathbf{Q}}_\perp^T \tilde{\mathbf{Q}}_\perp)^{-1} \tilde{\mathbf{Q}}_\perp^T \left[ \begin{pmatrix} \beta^k \\ \gamma^k \end{pmatrix} - \frac{1}{\rho} \begin{pmatrix} \mathbf{v}^{(3)k} \\ \mathbf{u}^{(2)k} \end{pmatrix} \right] \\ &= \tilde{\mathbf{Q}}_\perp \tilde{\mathbf{Q}}_\perp^T \left[ \begin{pmatrix} \beta^k \\ \gamma^k \end{pmatrix} - \frac{1}{\rho} \begin{pmatrix} \mathbf{v}^{(3)k} \\ \mathbf{u}^{(2)k} \end{pmatrix} \right] \\ &= (\mathbf{I}_{p+|\mathcal{T}|} - \tilde{\mathbf{Q}} \tilde{\mathbf{Q}}^T) \left[ \begin{pmatrix} \beta^k \\ \gamma^k \end{pmatrix} - \frac{1}{\rho} \begin{pmatrix} \mathbf{v}^{(3)k} \\ \mathbf{u}^{(2)k} \end{pmatrix} \right] \end{aligned}$$

5. Update global variables  $\beta$  and  $\gamma$ .

$$\beta^{k+1} := \arg \min_{\beta \in \mathbb{R}^p} \sum_{i=1}^3 \left\| \beta - \left( \beta^{(i)k+1} + \frac{1}{\rho} \mathbf{v}^{(i)k} \right) \right\|_2^2 = \bar{\beta}^{k+1} + \frac{1}{\rho} \bar{\mathbf{v}}^k \quad (11)$$

$$\gamma^{k+1} := \arg \min_{\gamma \in \mathbb{R}^{|\mathcal{T}|}} \sum_{j=1}^2 \left\| \gamma - \left( \gamma^{(j)k+1} + \frac{1}{\rho} \mathbf{u}^{(j)k} \right) \right\|_2^2 = \bar{\gamma}^{k+1} + \frac{1}{\rho} \bar{\mathbf{u}}^k \quad (12)$$

where  $\bar{\beta}^k := \frac{\beta^{(1)k} + \beta^{(2)k} + \beta^{(3)k}}{3}$ ,  $\bar{\mathbf{v}}^k := \frac{\mathbf{v}^{(1)k} + \mathbf{v}^{(2)k} + \mathbf{v}^{(3)k}}{3}$ ,  $\bar{\gamma}^k := \frac{\gamma^{(1)k} + \gamma^{(2)k}}{2}$  and  $\bar{\mathbf{u}}^k := \frac{\mathbf{u}^{(1)k} + \mathbf{u}^{(2)k}}{2}$ .

6. Update dual variables.

$$\begin{aligned} \mathbf{v}^{(1)k+1} &:= \mathbf{v}^{(i)k} + \rho(\beta^{(i)k+1} - \beta^{k+1}) \quad \text{for } i = 1, 2, 3, \\ \mathbf{u}^{(1)k+1} &:= \mathbf{u}^{(j)k} + \rho(\gamma^{(j)k+1} - \gamma^{k+1}) \quad \text{for } j = 1, 2. \end{aligned}$$

Similarly, averaging the updates for  $u$  and the updates for  $v$  gives

$$\bar{\mathbf{v}}^{k+1} = \bar{\mathbf{v}}^k + \rho(\bar{\beta}^{k+1} - \beta^{k+1}) \quad (13)$$

$$\bar{\mathbf{u}}^{k+1} = \bar{\mathbf{u}}^k + \rho(\bar{\gamma}^{k+1} - \gamma^{k+1}) \quad (14)$$

Substituting (11) and (12) into (13) and (14) yields that  $\bar{\mathbf{v}}^{k+1} = \bar{\mathbf{u}}^{k+1} = 0$  after the first iteration.

Using  $\beta^k = \bar{\beta}^k$  and  $\gamma^k = \bar{\gamma}^k$  in the above updates, the updates become Lines 9-16 of Algorithm 1. Next, we follow Section 3.3.1 in Boyd et al. (2011) to determine the termination criteria. We first write Problem (6) in the same form as Problem (3.1) in Boyd et al. (2011) which is presented below in typewriter font:

$$\min_{\mathbf{x}, \mathbf{z}} \{ \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{z}) \text{ s.t. } \mathbf{Ax} + \mathbf{Bz} = \mathbf{c} \}$$

where

$$\mathbf{A} = \mathbf{I}_{3p+2|\mathcal{T}|}, \mathbf{B} = - \begin{pmatrix} \mathbf{I}_p & 0 \\ \mathbf{I}_p & 0 \\ \mathbf{I}_p & 0 \\ 0 & \mathbf{I}_{|\mathcal{T}|} \\ 0 & \mathbf{I}_{|\mathcal{T}|} \end{pmatrix}, \mathbf{c} = 0, \mathbf{x} = \begin{pmatrix} \beta^{(1)} \\ \beta^{(2)} \\ \beta^{(3)} \\ \gamma^{(1)} \\ \gamma^{(2)} \end{pmatrix} \text{ and } \mathbf{z} = \begin{pmatrix} \beta \\ \gamma \end{pmatrix}.$$

The primal and dual residuals are

$$\mathbf{r}^{k+1} = \mathbf{Ax}^{k+1} + \mathbf{Bz}^{k+1} - \mathbf{c} = \begin{pmatrix} \beta^{(1)k+1} - \beta^{k+1} \\ \beta^{(2)k+1} - \beta^{k+1} \\ \beta^{(3)k+1} - \beta^{k+1} \\ \gamma^{(1)k+1} - \gamma^{k+1} \\ \gamma^{(2)k+1} - \gamma^{k+1} \end{pmatrix} \text{ and } \mathbf{s}^{k+1} = \rho \mathbf{A}^T \mathbf{B} (\mathbf{z}^{k+1} - \mathbf{z}^k) = \rho \begin{pmatrix} \beta^{k+1} - \beta^k \\ \beta^{k+1} - \beta^k \\ \beta^{k+1} - \beta^k \\ \gamma^{k+1} - \gamma^k \\ \gamma^{k+1} - \gamma^k \end{pmatrix}.$$

By Condition (3.12) in Boyd et al. (2011), the ADMM algorithm stops when both residuals are small. In our case, the termination criteria are the following.

1. The primal residual is small:

$$\begin{aligned} & \sqrt{\sum_{i=1}^3 \|\boldsymbol{\beta}^{(i)k} - \boldsymbol{\beta}^k\|_2^2 + \sum_{j=1}^2 \|\boldsymbol{\gamma}^{(j)k} - \boldsymbol{\gamma}^k\|_2^2} \\ & \leq \sqrt{3p + 2|\mathcal{T}|} \cdot \epsilon^{abs} + \epsilon^{rel} \cdot \max \left\{ \sqrt{\sum_{i=1}^3 \|\boldsymbol{\beta}^{(i)k}\|_2^2 + \sum_{j=1}^2 \|\boldsymbol{\gamma}^{(j)k}\|_2^2}, \sqrt{3 \|\boldsymbol{\beta}^k\|_2^2 + 2 \|\boldsymbol{\gamma}^k\|_2^2} \right\}. \end{aligned}$$

2. The dual residual is small:

$$\rho \cdot \sqrt{3\|\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k-1}\|_2^2 + 2\|\boldsymbol{\gamma}^k - \boldsymbol{\gamma}^{k-1}\|_2^2} \leq \sqrt{3p + 2|\mathcal{T}|} \cdot \epsilon^{abs} + \epsilon^{rel} \cdot \sqrt{\sum_{i=1}^3 \|\boldsymbol{v}^{(i)k}\|_2^2 + \sum_{j=1}^2 \|\boldsymbol{u}^{(j)k}\|_2^2}.$$

## D.2 Treatment of Intercept in Problem (5)

When an intercept  $\beta_0$  is included in the least squares, Problem (5) becomes:

$$\min_{\substack{\beta_0 \in \mathbb{R}, \boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\gamma} \in \mathbb{R}^{|\mathcal{T}|} \\ \text{s.t. } \boldsymbol{\beta} = \mathbf{A}\boldsymbol{\gamma}}} \left\{ \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \beta_0 \mathbf{1}_n\|_2^2 + \lambda \left( \alpha \sum_{\ell=1}^{|\mathcal{T}|} w_\ell |\gamma_\ell| + (1 - \alpha) \sum_{j=1}^p \tilde{w}_j |\beta_j| \right) \right\}. \quad (15)$$

First-order condition of the solution  $(\hat{\beta}_0, \hat{\boldsymbol{\beta}})$  yields that

$$\frac{\partial \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \beta_0 \mathbf{1}_n\|_2^2}{\partial \beta_0} \Big|_{(\beta_0, \boldsymbol{\beta}) = (\hat{\beta}_0, \hat{\boldsymbol{\beta}})} = \frac{1}{n} \mathbf{1}_n^T (\mathbf{1}_n \hat{\beta}_0 - (\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})) = \frac{1}{n} (n\hat{\beta}_0 - \mathbf{1}_n^T (\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})) = 0.$$

So  $\hat{\beta}_0 = \frac{1}{n} \mathbf{1}_n^T (\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})$ . Plugging  $\hat{\beta}_0$  in Problem (15) and letting  $\mathbf{H} = \mathbf{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T$  yields

$$\min_{\substack{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\gamma} \in \mathbb{R}^{|\mathcal{T}|} \\ \text{s.t. } \boldsymbol{\beta} = \mathbf{A}\boldsymbol{\gamma}}} \left\{ \frac{1}{2n} \|\mathbf{H}\mathbf{y} - \mathbf{H}\mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \left( \alpha \sum_{\ell=1}^{|\mathcal{T}|} w_\ell |\gamma_\ell| + (1 - \alpha) \sum_{j=1}^p \tilde{w}_j |\beta_j| \right) \right\},$$

which can now be solved using our consensus ADMM algorithm.

## E Proof of Lemma 1

We first show existence of such  $B^*$  by providing a feasible procedure to find  $B^*$ . Suppose  $\boldsymbol{\beta}^*$  has at least two distinct values (otherwise  $B^* = \{r\}$  trivially). Start with  $B = \mathcal{L}(\mathcal{T})$  so that the first constraint is satisfied. If for siblings  $u, v$  in  $B$  such that the second constraint is violated, by construction  $\beta_j^* = \beta_k^*$  for  $j \in \mathcal{L}(\mathcal{T}_u)$  and  $k \in \mathcal{L}(\mathcal{T}_v)$ . So we replace  $u, v$  in  $B$  with their parent node. We repeat the above steps until the second constraint is satisfied, while holding the first constraint. Thus,  $B$  satisfies the two requirements for  $B^*$ .

Suppose  $B^*$  and  $\tilde{B}^*$  are different aggregating sets for  $\boldsymbol{\beta}^*$ . Without loss of generality, suppose there exists  $u \in \tilde{B}^*$  but  $u \notin B^*$ . Then  $u$  is a descendant or an ancestor of some nodes in  $B^*$ ; for either case the second constraint will be violated. Thus, such  $u$  does not exist and  $\tilde{B}^* = B^*$ .

The existence and uniqueness of  $\mathcal{A}^*$  follow from the definition of support of  $\boldsymbol{\beta}^*$ .

## F Proof of Theorem 4

Denote our penalty by  $\Omega(\boldsymbol{\beta}, \boldsymbol{\gamma}) := \alpha \sum_{\ell=1}^{|T|} w_\ell |\gamma_\ell| + (1 - \alpha) \sum_{j=1}^p \tilde{w}_j |\beta_j|$ . We follow the proof strategy used in Theorem 1 of Lou et al. (2016) to prove this theorem. If  $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$  is a solution to Problem (5), then we have

$$\frac{1}{2n} \left\| \mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}} \right\|_2^2 + \lambda \Omega(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) \leq \frac{1}{2n} \left\| \mathbf{y} - \mathbf{X} \boldsymbol{\beta} \right\|_2^2 + \lambda \Omega(\boldsymbol{\beta}, \boldsymbol{\gamma})$$

for any  $(\boldsymbol{\beta}, \boldsymbol{\gamma})$  such that  $\boldsymbol{\beta} = \mathbf{A} \boldsymbol{\gamma}$ . Let  $(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*)$  be such that

$$\boldsymbol{\beta}^* = \mathbf{A}_{B^*} \tilde{\boldsymbol{\beta}}^* \quad \text{and} \quad \gamma_\ell^* = \begin{cases} \tilde{\beta}_\ell^* & \text{if } \ell \in B^* \\ 0 & \text{otherwise.} \end{cases}$$

Plugging in  $\mathbf{y} = \mathbf{X} \boldsymbol{\beta}^* + \boldsymbol{\varepsilon}$  and  $(\boldsymbol{\beta}, \boldsymbol{\gamma}) = (\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*)$ , with some algebra we have

$$\frac{1}{2n} \left\| \mathbf{X} \hat{\boldsymbol{\beta}} - \mathbf{X} \boldsymbol{\beta}^* \right\|_2^2 + \lambda \Omega(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) \leq \lambda \Omega(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) + \frac{1}{n} \boldsymbol{\varepsilon}^T \mathbf{X} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)} \quad (16)$$

where  $\hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)} = \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*$ . By  $\hat{\boldsymbol{\beta}} = \mathbf{A} \hat{\boldsymbol{\gamma}}$  and  $\boldsymbol{\beta}^* = \mathbf{A} \boldsymbol{\gamma}^*$  (and writing  $\hat{\boldsymbol{\Delta}}^{(\boldsymbol{\gamma}^*)} = \hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma}^*$ ),

$$\frac{1}{n} \boldsymbol{\varepsilon}^T \mathbf{X} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)} = \frac{1}{n} \boldsymbol{\varepsilon}^T \mathbf{X} \mathbf{A} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\gamma}^*)}.$$

We next bound  $n^{-1} \boldsymbol{\varepsilon}^T \mathbf{X} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)} = (1 - \alpha) n^{-1} \boldsymbol{\varepsilon}^T \mathbf{X} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)} + \alpha n^{-1} \boldsymbol{\varepsilon}^T \mathbf{X} \mathbf{A} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\gamma}^*)}$  in absolute value. Define  $V_j := n^{-1/2} \mathbf{X}_j^T \boldsymbol{\varepsilon}$  for  $j = 1, \dots, p$  and  $U_\ell := n^{-1/2} \mathbf{A}_\ell^T \mathbf{X}^T \boldsymbol{\varepsilon}$  for  $\ell = 1, \dots, |T|$ . With the choice of weights  $\tilde{w}_j = \|\mathbf{X}_j\|_2 / \sqrt{n}$  for  $1 \leq j \leq p$  and  $w_\ell = \|\mathbf{X} \mathbf{A}_\ell\|_2 / \sqrt{n}$  for  $1 \leq \ell \leq |T|$ ,

$$\begin{aligned} \left| \frac{1}{n} \boldsymbol{\varepsilon}^T \mathbf{X} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)} \right| &= \left| (1 - \alpha) \frac{1}{n} \boldsymbol{\varepsilon}^T \mathbf{X} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)} + \alpha \frac{1}{n} \boldsymbol{\varepsilon}^T \mathbf{X} \mathbf{A} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\gamma}^*)} \right| \\ &= \left| (1 - \alpha) \sum_{j=1}^p \left\{ \frac{\mathbf{X}_j^T \boldsymbol{\varepsilon}}{\sqrt{n} \|\mathbf{X}_j\|_2} \cdot \frac{\|\mathbf{X}_j\|_2}{\sqrt{n}} \hat{\Delta}_j^{(\boldsymbol{\beta}^*)} \right\} + \alpha \sum_{\ell=1}^{|T|} \left\{ \frac{\mathbf{A}_\ell^T \mathbf{X}^T \boldsymbol{\varepsilon}}{\sqrt{n} \|\mathbf{X} \mathbf{A}_\ell\|_2} \cdot \frac{\|\mathbf{X} \mathbf{A}_\ell\|_2}{\sqrt{n}} \hat{\Delta}_\ell^{(\boldsymbol{\gamma}^*)} \right\} \right| \\ &= \left| (1 - \alpha) \sum_{j=1}^p \left\{ \frac{V_j}{\|\mathbf{X}_j\|_2} \cdot \tilde{w}_j \hat{\Delta}_j^{(\boldsymbol{\beta}^*)} \right\} + \alpha \sum_{\ell=1}^{|T|} \left\{ \frac{U_\ell}{\|\mathbf{X} \mathbf{A}_\ell\|_2} \cdot w_\ell \hat{\Delta}_\ell^{(\boldsymbol{\gamma}^*)} \right\} \right| \\ &\leq (1 - \alpha) \sum_{j=1}^p \left\{ \frac{|V_j|}{\|\mathbf{X}_j\|_2} \cdot \tilde{w}_j |\hat{\Delta}_j^{(\boldsymbol{\beta}^*)}| \right\} + \alpha \sum_{\ell=1}^{|T|} \left\{ \frac{|U_\ell|}{\|\mathbf{X} \mathbf{A}_\ell\|_2} \cdot w_\ell |\hat{\Delta}_\ell^{(\boldsymbol{\gamma}^*)}| \right\} \\ &\leq (1 - \alpha) \max_{1 \leq j \leq p} \frac{|V_j|}{\|\mathbf{X}_j\|_2} \cdot \sum_{j=1}^p \tilde{w}_j |\hat{\Delta}_j^{(\boldsymbol{\beta}^*)}| + \alpha \max_{1 \leq \ell \leq |T|} \frac{|U_\ell|}{\|\mathbf{X} \mathbf{A}_\ell\|_2} \cdot \sum_{\ell=1}^{|T|} w_\ell |\hat{\Delta}_\ell^{(\boldsymbol{\gamma}^*)}| \end{aligned} \quad (17)$$

where the second to last inequality follows from the triangle inequality.

Since  $\boldsymbol{\varepsilon} \sim N_n(0, \sigma^2 \mathbf{I}_n)$ ,  $V_j \sim N(0, \|\mathbf{X}_j\|_2^2 \sigma^2/n)$  for  $j = 1, \dots, p$  and  $U_\ell \sim N(0, \|\mathbf{X} \mathbf{A}_\ell\|_2^2 \sigma^2/n)$  for  $\ell = 1, \dots, |\mathcal{T}|$ . By Lemma 6.2 of Bühlmann and van de Geer (2011), we have for  $x > 0$

$$\begin{aligned} \mathbb{P} \left( \max_{1 \leq j \leq p} \frac{|V_j|}{\|\mathbf{X}_j\|_2} > \sigma \sqrt{\frac{2(x + \log p)}{n}} \right) &\leq 2e^{-x} \quad \text{and} \\ \mathbb{P} \left( \max_{1 \leq \ell \leq |\mathcal{T}|} \frac{|U_\ell|}{\|\mathbf{X} \mathbf{A}_\ell\|_2} > \sigma \sqrt{\frac{2(x + \log |\mathcal{T}|)}{n}} \right) &\leq 2e^{-x}. \end{aligned}$$

By the union bound,

$$\mathbb{P} \left( \max_{1 \leq j \leq p} \frac{|V_j|}{\|\mathbf{X}_j\|_2} > \sigma \sqrt{\frac{2(x + \log p)}{n}} \quad \text{or} \quad \max_{1 \leq \ell \leq |\mathcal{T}|} \frac{|U_\ell|}{\|\mathbf{X} \mathbf{A}_\ell\|_2} > \sigma \sqrt{\frac{2(x + \log |\mathcal{T}|)}{n}} \right) \leq 4e^{-x}.$$

By the construction of  $\mathcal{T}$ , each internal node has at least 2 child nodes. To go up to the next level from the leaf nodes, only one node “survives” among its siblings. For  $\mathcal{T}$  with  $p$  leaf nodes, there must be at most  $p - 1$  internal nodes where the maximum number is achieved when  $\mathcal{T}$  is a full binary tree. So  $|\mathcal{T}| \leq 2p$ . Choosing  $x = \log(2p)$ , we have that the following results hold with probability at least  $1 - 2/p$ :

$$\max_{1 \leq j \leq p} \frac{|V_j|}{\|\mathbf{X}_j\|_2} \leq 2\sigma \sqrt{\frac{\log(2p)}{n}} \quad \text{and} \quad \max_{1 \leq \ell \leq |\mathcal{T}|} \frac{|U_\ell|}{\|\mathbf{X} \mathbf{A}_\ell\|_2} \leq 2\sigma \sqrt{\frac{\log(2p)}{n}}.$$

Plugging these upper bounds into (17), we have the following inequality holding with high probability:

$$\begin{aligned} \left| \frac{1}{n} \boldsymbol{\varepsilon}^T \mathbf{X} \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)} \right| &\leq 2\sigma \sqrt{\frac{\log(2p)}{n}} \left( (1 - \alpha) \sum_{j=1}^p \tilde{w}_j \left| \hat{\Delta}_j^{(\boldsymbol{\beta}^*)} \right| + \alpha \sum_{\ell=1}^{|\mathcal{T}|} w_\ell \left| \hat{\Delta}_\ell^{(\boldsymbol{\gamma}^*)} \right| \right) \\ &= 2\sigma \sqrt{\frac{\log(2p)}{n}} \Omega(\hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)}, \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\gamma}^*)}) \end{aligned} \tag{18}$$

Let  $\lambda \geq 4\sigma \sqrt{\log(2p)/n}$  and  $0 \leq \alpha \leq 1$ . By (16) and (18), the following holds with probability at least  $1 - 2/p$ :

$$\begin{aligned} \frac{1}{2n} \left\| \mathbf{X} \hat{\boldsymbol{\beta}} - \mathbf{X} \boldsymbol{\beta}^* \right\|_2^2 &\leq \frac{1}{2} \lambda \Omega(\hat{\boldsymbol{\Delta}}^{(\boldsymbol{\beta}^*)}, \hat{\boldsymbol{\Delta}}^{(\boldsymbol{\gamma}^*)}) - \lambda \Omega(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) + \lambda \Omega(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) \\ &\leq \frac{1}{2} \left( \lambda \Omega(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) + \lambda \Omega(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) \right) - \lambda \Omega(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) + \lambda \Omega(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) \quad (\text{by triangle inequality}) \\ &\leq \frac{3}{2} \lambda \Omega(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) = \frac{3}{2} \lambda \left( (1 - \alpha) \sum_{j=1}^p \tilde{w}_j |\beta_j^*| + \alpha \sum_{\ell=1}^{|\mathcal{T}|} w_\ell |\gamma_\ell^*| \right) \\ &= \frac{3}{2} \lambda \left( (1 - \alpha) \sum_{j \in \mathcal{A}^*} \tilde{w}_j |\beta_j^*| + \alpha \sum_{\ell \in B^*} w_\ell |\tilde{\beta}_\ell^*| \right) \end{aligned}$$

where  $\mathcal{A}^*$  is the support of  $\boldsymbol{\beta}^*$ .

## G Proof of Corollary 1

With  $\|\mathbf{X}_j\|_2 \leq \sqrt{n}$  the weights in Theorem 4 become:

- For  $1 \leq j \leq p$ ,  $\tilde{w}_j = \frac{\|\mathbf{X}_j\|_2}{\sqrt{n}} \leq 1$ .
- For  $1 \leq \ell \leq |\mathcal{T}|$ ,  $w_\ell = \frac{\|\mathbf{X}\mathbf{A}_\ell\|_2}{\sqrt{n}} \leq \frac{1}{\sqrt{n}} \sum_{j \in \mathcal{L}(\mathcal{T}_{u_\ell})} \|\mathbf{X}_j\|_2 \leq |\mathcal{L}(\mathcal{T}_{u_\ell})|$ .

The inequality in  $w_\ell$  is by the triangle inequality and the definition of  $\mathbf{A}$ : the  $\ell$ th column  $\mathbf{A}_\ell$  encodes the descendant leaves of the node  $u_\ell$ , i.e.,  $A_{j\ell} = 1_{\{j \in \text{descendant}(u_\ell) \cup \{u_\ell\}\}}$ . By construction of  $\tilde{\beta}^*$ , all coefficients in the branch rooted by an aggregating node  $u_\ell$  share the same value, i.e.,  $\beta_j^* = \tilde{\beta}_\ell^*$  for all  $j \in \mathcal{L}(\mathcal{T}_{u_\ell})$  where  $\ell \in B^*$ . Thus, for  $\ell \in B^*$ ,

$$\tilde{\beta}_\ell^* = \frac{\|\beta_{\mathcal{L}(\mathcal{T}_{u_\ell})}^*\|_1}{|\mathcal{L}(\mathcal{T}_{u_\ell})|} \Rightarrow \sum_{\ell \in B^*} w_\ell |\tilde{\beta}_\ell^*| \leq \sum_{\ell \in B^*} |\mathcal{L}(\mathcal{T}_{u_\ell})| \left( \frac{\|\beta_{\mathcal{L}(\mathcal{T}_{u_\ell})}^*\|_1}{|\mathcal{L}(\mathcal{T}_{u_\ell})|} \right) = \|\beta^*\|_1.$$

Plugging the choice of  $\lambda$  and the above upper bounds into Theorem (4) yields

$$\begin{aligned} \frac{1}{n} \|\mathbf{X}\hat{\beta} - \mathbf{X}\beta^*\|_2^2 &\leq 12\sigma \sqrt{\frac{\log(2p)}{n}} \left( (1-\alpha) \sum_{j \in \mathcal{A}^*} \tilde{w}_j |\beta_j^*| + \alpha \sum_{\ell \in B^*} w_\ell |\tilde{\beta}_\ell^*| \right) \\ &\leq 12\sigma \sqrt{\frac{\log(2p)}{n}} ((1-\alpha) \|\beta^*\|_1 + \alpha \|\beta^*\|_1) \\ &\lesssim \sigma \sqrt{\frac{\log p}{n}} \|\beta^*\|_1 \quad (\text{by } \log(2p) = \log 2 + \log p \leq 2 \log p) \end{aligned}$$

which holds up to a multiplicative constant with probability at least  $1 - 2/p$ .

## H Proof of Corollary 2

With  $\beta_1^* = \dots = \beta_p^*$  the aggregating set consists of the root, i.e.,  $B^* = \{r\}$ . Thus,  $\tilde{\beta}_r^* = \beta_1^* = \|\beta^*\|_1/p$  and  $w_r = \|\mathbf{X}\mathbf{1}_p\|_2/\sqrt{n}$ . With  $\|\mathbf{X}_j\|_2 \leq \sqrt{n}$ , we have  $\tilde{w}_j = \|\mathbf{X}_j\|_2/\sqrt{n} \leq 1$  for  $1 \leq j \leq p$ .

Plugging  $\lambda = 4\sigma\sqrt{\log(2p)/n}$  into Theorem (4) yields

$$\begin{aligned} \frac{1}{n} \|\mathbf{X}\hat{\beta} - \mathbf{X}\beta^*\|_2^2 &\leq 12\sigma \sqrt{\frac{\log(2p)}{n}} \left( (1-\alpha) \sum_{j \in \mathcal{A}^*} \tilde{w}_j |\beta_j^*| + \alpha w_r |\tilde{\beta}_r^*| \right) \\ &\leq 12\sigma \sqrt{\frac{\log(2p)}{n}} \left( (1-\alpha) \|\beta^*\|_1 + \alpha \frac{\|\mathbf{X}\mathbf{1}_p\|_1}{p\sqrt{n}} \|\beta^*\|_1 \right) \\ &\lesssim \sigma \sqrt{\frac{\log p}{n}} \|\beta^*\|_1 \left( (1-\alpha) + \alpha \frac{\|\mathbf{X}\mathbf{1}_p\|_2}{p\sqrt{n}} \right) \end{aligned}$$

which holds up to a multiplicative constant with probability at least  $1 - 2/p$ . Now, if  $\alpha \geq [1 + \|\mathbf{X}\mathbf{1}_p\|_2/(p\sqrt{n})]^{-1}$ , then

$$1 - \alpha \leq \alpha \frac{\|\mathbf{X}\mathbf{1}_p\|_2}{p\sqrt{n}}$$

and so

$$\frac{1}{n} \|\mathbf{X}\hat{\boldsymbol{\beta}} - \mathbf{X}\boldsymbol{\beta}^*\|_2^2 \lesssim \sigma \sqrt{\frac{\log p}{n}} \|\boldsymbol{\beta}^*\|_1 \alpha \frac{\|\mathbf{X}\mathbf{1}_p\|_2}{p\sqrt{n}}$$

It follows that if  $\|\mathbf{X}\mathbf{1}_p\|_2 = o(p\sqrt{n})$ , then our method achieves a better rate than the lasso rate of  $\sigma\sqrt{\log p/n} \|\boldsymbol{\beta}^*\|_1$ .

## I Proof of Proposition 1

Let  $\dot{\mathbf{X}}^{(1)}, \dots, \dot{\mathbf{X}}^{(p)} \in \{0, 1\}^n$  be *iid* copies of the random vector  $\dot{\mathbf{X}} \in \{0, 1\}^n$  with  $\mathbf{1}_n^T \dot{\mathbf{X}} = r$  and  $\mathbb{P}(\dot{X}_i = 1) = r/n \ \forall i = 1, \dots, n$ . By construction  $\{\dot{\mathbf{X}}^{(1)}, \dots, \dot{\mathbf{X}}^{(p)}\}$  each has  $r$  nonzero entries. For the  $j$ th column of  $\mathbf{X}$ , let  $\mathbf{X}_j = \sqrt{n/r} \dot{\mathbf{X}}^{(j)}$  so that its nonzero entries equal  $\sqrt{n/r}$  and  $\|\mathbf{X}_j\|_2 = \sqrt{n}$ .

For the random vector  $\dot{\mathbf{X}}$ , its expectation and covariance ( $\boldsymbol{\Sigma}$ ) diagonal entries are

$$\mathbb{E}\dot{\mathbf{X}} = \frac{r}{n}\mathbf{1}_n \quad \text{and} \quad \Sigma_{ii} = \text{Var}(\dot{X}_i) = \frac{r}{n} \left(1 - \frac{r}{n}\right) =: \tau^2.$$

Suppose  $\|\dot{\mathbf{X}} - \mathbb{E}\dot{\mathbf{X}}\|_2 \leq L$  for some  $L > 0$ . Let  $\mathbf{Z} := \sum_{j=1}^p \dot{\mathbf{X}}^{(j)}$  and  $v(\mathbf{Z}) := \max\{\|\mathbb{E}[(\mathbf{Z} - \mathbb{E}\mathbf{Z})(\mathbf{Z} - \mathbb{E}\mathbf{Z})^T]\|_{op}, \mathbb{E}[\|\mathbf{Z} - \mathbb{E}\mathbf{Z}\|_2^2]\}$ . Corollary 6.1.2 of Tropp (2015) gives the concentration inequality for  $\mathbf{Z}$ :  $\forall t \geq 0$ ,

$$\mathbb{P}(\|\mathbf{Z} - \mathbb{E}\mathbf{Z}\|_2 \geq t) \leq (n+1) \cdot \exp\left(\frac{-t^2/2}{v(\mathbf{Z}) + Lt/3}\right). \quad (19)$$

First, we derive a closed form for  $L$ . By construction of  $\dot{\mathbf{X}}$  (in particular  $\mathbf{1}_n^T \dot{\mathbf{X}} = r$ ),

$$\left\|\dot{\mathbf{X}} - \mathbb{E}\dot{\mathbf{X}}\right\|_2^2 = \left\|\dot{\mathbf{X}} - \frac{r}{n}\mathbf{1}_n\right\|_2^2 = (n-r)(0 - \frac{r}{n})^2 + r(1 - \frac{r}{n})^2 = r(1 - \frac{r}{n}) = n\tau^2.$$

So choose  $L = \|\dot{\mathbf{X}} - \mathbb{E}\dot{\mathbf{X}}\|_2 = \sqrt{n}\tau$ . Next, we simplify the two components of  $v(\mathbf{Z})$ :

$$\|\mathbb{E}[(\mathbf{Z} - \mathbb{E}\mathbf{Z})(\mathbf{Z} - \mathbb{E}\mathbf{Z})^T]\|_{op} = \|\text{Cov}(\mathbf{Z})\|_{op} = p\|\boldsymbol{\Sigma}\|_{op},$$

$$\mathbb{E}[\|\mathbf{Z} - \mathbb{E}\mathbf{Z}\|_2^2] = \sum_{i=1}^n \text{Var}(Z_i) = \text{tr}(\text{Cov}(\mathbf{Z})) = p \cdot \text{tr}(\boldsymbol{\Sigma}).$$

Furthermore,  $\max\{\|\boldsymbol{\Sigma}\|_{op}, \text{tr}(\boldsymbol{\Sigma})\} = \text{tr}(\boldsymbol{\Sigma})$  since  $\boldsymbol{\Sigma}$  is positive semidefinite. Combining the above yields

$$v(\mathbf{Z}) = p \cdot \text{tr}(\boldsymbol{\Sigma}) = p \cdot n \cdot \tau^2 = pL^2.$$

Substituting  $v(\mathbf{Z})$  and  $\mathbb{E}\mathbf{Z} = pr/n \cdot \mathbf{1}_n$  into (19) yields

$$\mathbb{P}\left(\left\|\mathbf{Z} - \frac{pr}{n}\mathbf{1}_n\right\|_2 \geq t\right) \leq (n+1) \cdot \exp\left(\frac{-t^2/2}{pL^2 + Lt/3}\right) \quad \forall t \geq 0.$$

Since  $\mathbf{X}_j = \sqrt{n/r} \hat{\mathbf{X}}^{(j)}$ , we have  $\mathbf{X} \mathbf{1}_p = \sqrt{n/r} \mathbf{Z}$  and  $\forall t \geq 0$ ,

$$\mathbb{P} \left( \left\| \mathbf{X} \mathbf{1}_p - p \sqrt{\frac{r}{n}} \mathbf{1}_n \right\|_2 \geq \sqrt{\frac{n}{r}} t \right) = \mathbb{P} \left( \left\| \mathbf{Z} - \frac{pr}{n} \mathbf{1}_n \right\|_2 \geq t \right) \leq (n+1) \cdot \exp \left( \frac{-(t/L)^2/2}{p + (t/L)/3} \right) =: F(t).$$

Choosing  $t = pr/\sqrt{n}$  and by the triangle inequality,

$$\|\mathbf{X} \mathbf{1}_p\|_2 \leq \|\mathbf{X} \mathbf{1}_p - p \sqrt{r/n} \cdot \mathbf{1}_n\|_2 + \|p \sqrt{r/n} \cdot \mathbf{1}_n\|_2 \leq \sqrt{n/r} \cdot t + p \sqrt{r} = 2p \sqrt{r}$$

holds with probability at least  $1 - F(pr/\sqrt{n})$ . Also  $t/L = pr/(n\tau)$ . We next simplify  $F(pr/\sqrt{n})$ .

$$\begin{aligned} F \left( \frac{pr}{\sqrt{n}} \right) &= (n+1) \cdot \exp \left( \frac{-\frac{1}{2} \left( \frac{pr}{n\tau} \right)^2}{p + \frac{1}{3} \left( \frac{pr}{n\tau} \right)} \right) = (n+1) \cdot \exp \left( \frac{-\frac{pr^2}{2n^2\tau^2}}{1 + \frac{r}{3n\tau}} \right) \\ \Rightarrow \log F \left( \frac{pr}{\sqrt{n}} \right) &= \log(n+1) - \frac{1}{2} \cdot \frac{p(r/n)^2}{\tau^2 + \frac{1}{3} \cdot \tau \cdot \frac{r}{n}} \\ &\leq \log(n+1) - \frac{1}{2} \cdot \frac{p(r/n)^2}{r/n + \frac{1}{3}(r/n)^{3/2}} \quad \left( \text{by } \tau = \sqrt{\frac{r}{n} \left( 1 - \frac{r}{n} \right)} \leq \sqrt{\frac{r}{n}} \right) \\ &\leq \log(n+1) - \frac{1}{2} \cdot \frac{p(r/n)}{1 + \frac{1}{3}} \quad \left( \text{by } \frac{r}{n} \leq 1 \right) \\ &= \log(n+1) - \frac{3}{8} \cdot \frac{pr}{n}. \end{aligned}$$

So,  $\mathbb{P}(\|\mathbf{X} \mathbf{1}_p\|_2 \leq 2p \sqrt{r}) \geq 1 - \exp \left( \log(n+1) - \frac{3pr}{8n} \right)$ . Note that  $\frac{r}{n} > \frac{8 \log(n+1)}{3p}$  ensures the exponent converges to negative infinity as  $n$  and  $p$  increase. Thus, if  $r > \frac{8n \log(n+1)}{3p}$ , then

$$\|\mathbf{X} \mathbf{1}_p\|_2 = O_p(p \sqrt{r}) \Leftrightarrow \frac{\|\mathbf{X} \mathbf{1}_p\|_2}{p \sqrt{n}} = O_p \left( \sqrt{\frac{r}{n}} \right).$$

So, we have  $\|\mathbf{X} \mathbf{1}_p\|_2 = o_p(p \sqrt{n})$  if  $r = o(n)$  and  $r > \frac{8n \log(n+1)}{3p}$ .

## J Paired $t$ -tests between our method and L1-ag-h

To compare the test errors in the first two columns of Table 1 in the main paper, we perform a two-sample  $t$ -test for each row. For each of the  $n = 40,000$  reviews, a pair of squared prediction errors is calculated, one from each method. Conditional on the training set, these  $n$  pairs can be thought of as independent. Table 4 shows the resulting p-values.

prop.	1%	5%	10%	20%	40%	60%	80%	100%
p-value	$3.95 \times 10^{-9}$	0.086	$7.21 \times 10^{-7}$	$7.56 \times 10^{-8}$	$2.2 \times 10^{-16}$	0.056	0.105	0.001

Table 4: Paired  $t$ -tests between squared prediction errors from our method and L1-ag-h on the held-out test set. The prop. values correspond to different training sizes in Table 1.

## K Hierarchical Clustering Tree with ELMo

Embeddings from Language Models (ELMo) is a state-of-the-art natural language processing framework for representing words with deep contextualized embeddings (Peters et al., 2018). ELMo leverages a deep neural network that is pre-trained on a large text corpus and has multiple attractive properties. For our purpose, the most relevant property of ELMo is that its word embeddings account for the *context* in which the words appear. In particular, the ELMo embedding assigned to a word is a function of the entire sentence containing that word. This means that the same word in different contexts will have different embeddings. This is a great advantage over traditional word embeddings such as *GloVe* (Pennington et al., 2014).

Among the training sets listed in Table 1, we focus on one scenario when  $n = 33,997$  and  $p = 5,621$  (which corresponds to 20% of the entire training reviews). We use TensorFlow (Abadi et al., 2015) and an ELMo model pre-trained on the One Billion World Language Model Benchmark<sup>5</sup> to get 1,024-dimensional embeddings for all words in the 33,997 training reviews and the 40,000 held-out testing reviews. To establish a fair comparison with existing experimentation, we keep only adjective embeddings in the reviews, which leaves us 952,974 adjective embeddings for training and 1,122,030 adjective embeddings for testing. Again, the same adjective in various sentences has different ELMo embeddings whose proximities encode how the word is used in each context.

We use mini-batch K-Means clustering (Sculley, 2010) to cluster the 952,974 adjectives from training into 6,001 clusters. The choice of the number of K-Means clusters ensures that the comparison between ELMo and *GloVe* (which is based on  $p=5,621$ ) is roughly on the same scale of features. After mini-batch K-Means clustering, we update each cluster centroid to the averaged embeddings for training adjectives assigned to the cluster. We also associate testing adjectives to the clusters by Euclidean distance. After associating adjectives to clusters, we construct a document-cluster matrix for training and testing, respectively. The  $ij$ -entry of the document-cluster matrix has the number occurrences of the  $j$ th cluster in the  $i$ th review. Each cluster represents a collection of occurrences of adjectives used in a similar context within a review. We next generate a tree by performing hierarchical clustering on the cluster centroids. The tree and document-cluster matrix is then used in fitting our model along with others such as L1-ag-h and L1-ag-d.

## References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org.

---

<sup>5</sup><http://www.statmt.org/lm-benchmark/>

Arnold, T. B. and Tibshirani, R. J. (2014). *genlasso: Path algorithm for generalized lasso problems*. R package version 1.3.

Bernstein, D. S. (2009). *Matrix Mathematics: Theory, Facts, and Formulas (Second Edition)*. Princeton University Press.

Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122.

Bühlmann, P. and van de Geer, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Publishing Company, Incorporated, 1st edition.

Cao, Y., Zhang, A., and Li, H. (2017). Microbial Composition Estimation from Sparse Count Data. *ArXiv e-prints*.

Caporaso, J. G., Kuczynski, J., Stombaugh, J., Bittinger, K., Bushman, F. D., Costello, E. K., Fierer, N., Pena, A. G., Goodrich, J. K., Gordon, J. I., Huttley, G. A., Kelley, S. T., Knights, D., Koenig, J. E., Ley, R. E., Lozupone, C. A., McDonald, D., Muegge, B. D., Pirrung, M., Reeder, J., Sevinsky, J. R., Turnbaugh, P. J., Walters, W. A., Widmann, J., Yatsunenko, T., Zaneveld, J., and Knight, R. (2010). QIIME allows analysis of high-throughput community sequencing data. *Nat. Methods*, 7(5):335–336.

Chen, J., Bushman, F. D., Lewis, J. D., Wu, G. D., and Li, H. (2013). Structure-constrained sparse canonical correlation analysis with an application to microbiome data analysis. *Biostatistics*, 14(2):244–258.

Côté, F. D., Psaromiligkos, I. N., and Gross, W. J. (2012). A Chernoff-type Lower Bound for the Gaussian Q-function. *ArXiv e-prints*.

Feinerer, I. and Hornik, K. (2016). *wordnet: WordNet Interface*. R package version 0.1-11.

Feinerer, I. and Hornik, K. (2017). *tm: Text Mining Package*. R package version 0.7-1.

Fellbaum, C. (1998). *WordNet: An Electronic Lexical Database*. Bradford Books.

Forman, G. (2003). An extensive empirical study of feature selection metrics for text classification. *J. Mach. Learn. Res.*, 3:1289–1305.

Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22.

Guinot, F., Szafranski, M., Ambroise, C., and Samson, F. (2017). Learning the optimal scale for GWAS through hierarchical SNP aggregation. *ArXiv e-prints*.

Huang, A. (2008). Similarity measures for text document clustering. In *Proceedings of the sixth new zealand computer science research student conference (NZCSRSC2008), Christchurch, New Zealand*, pages 49–56.

Ke, T., Fan, J., and Wu, Y. (2015). Homogeneity Pursuit. *J Am Stat Assoc*, 110(509):175–194.

Khabbazian, M., Kriebel, R., Rohe, K., and An, C. (2016). Fast and accurate detection of evolutionary shifts in ornsteinuhlenbeck models. *Methods in Ecology and Evolution*, 7(7):811–824.

Kim, S., Xing, E. P., et al. (2012). Tree-guided group lasso for multi-response regression with structured sparsity, with an application to eqtl mapping. *The Annals of Applied Statistics*, 6(3):1095–1117.

Li, C. and Li, H. (2010). Variable selection and regression analysis for graph-structured covariates with an application to genomics. *Ann. Appl. Stat.*, 4(3):1498–1516.

Li, Y., Raskutti, G., and Willett, R. (2018). Graph-based regularization for regression problems with highly-correlated designs. *ArXiv e-prints*.

Lin, W., Shi, P., Feng, R., and Li, H. (2014). Variable selection in regression with compositional covariates. *Biometrika*, 101:785–797.

Liu, X., Yu, S., Janssens, F., Gläzel, W., Moreau, Y., and De Moor, B. (2010). Weighted hybrid clustering by combining text mining and bibliometrics on a large-scale journal database. *J. Am. Soc. Inf. Sci. Technol.*, 61(6):1105–1119.

Lou, Y., Bien, J., Caruana, R., and Gehrke, J. (2016). Sparse partially linear additive models. *Journal of Computational and Graphical Statistics*, 25(4):1126–1140.

Matsen, F. A., Kodner, R. B., and Armbrust, E. V. (2010). pplacer: linear time maximum-likelihood and Bayesian phylogenetic placement of sequences onto a fixed reference tree. *BMC Bioinformatics*, 11:538.

McMurdie, P. J. and Holmes, S. (2013). phyloseq: An r package for reproducible interactive analysis and graphics of microbiome census data. *PLOS ONE*, 8(4):1–11.

Mohammad, S. M. and Turney, P. D. (2013). Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29(3):436–465.

Mukherjee, R., Pillai, N. S., and Lin, X. (2015). Hypothesis testing for high-dimensional sparse binary regression. *Ann. Statist.*, 43(1):352–381.

Pennington, J., Socher, R., and Manning, C. D. (2014). Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). Deep contextualized word representations. In *Proc. of NAACL*.

R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Randolph, T. W., Zhao, S., Copeland, W., Hullar, M., and Shojaie, A. (2015). Kernel-Penalized Regression for Analysis of Microbiome Data. *ArXiv e-prints*.

Ridenhour, B. J., Brooker, S. L., Williams, J. E., Van Leuven, J. T., Miller, A. W., Dearing, M. D., and Remien, C. H. (2017). Modeling time-series data from microbial communities. *ISME J*, 11(11):2526–2537.

Schloss, P., L Westcott, S., Ryabin, T., R Hall, J., Hartmann, M., Hollister, E., Lesniewski, R., Oakley, B., Parks, D., Robinson, C., W Sahl, J., Stres, B., G Thallinger, G., Van Horn, D., and Weber, C. (2009). Introducing mothur: Open-source, platform-independent, community-supported software for describing and comparing microbial communities. *Applied and environmental microbiology*, 75(23):7537–41.

Sculley, D. (2010). Web-scale k-means clustering. In *Proceedings of the 19th International Conference on World Wide Web*, WWW 10, page 11771178, New York, NY, USA. Association for Computing Machinery.

She, Y. (2010). Sparse regression with exact clustering. *Electron. J. Statist.*, 4:1055–1096.

Shi, P., Zhang, A., and Li, H. (2016). Regression analysis for microbiome compositional data. *Ann. Appl. Stat.*, 10(2):1019–1040.

Tang, Y., Li, M., and Niclaoe, D. L. (2016). Phylogenetic Dirichlet-multinomial model for microbiome data. *ArXiv e-prints*.

Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., and Kappas, A. (2010). Sentiment in short strength detection informal text. *J. Am. Soc. Inf. Sci. Technol.*, 61(12):2544–2558.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288.

Tibshirani, R. J. and Taylor, J. (2011). The solution path of the generalized lasso. *Ann. Statist.*, 39(3):1335–1371.

Tropp, J. A. (2015). An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230.

Wallace, M. (2007). *Jawbone Java WordNet API*.

Wang, H., Lu, Y., and Zhai, C. (2010). Latent aspect rating analysis on review text data: A rating regression approach. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '10, pages 783–792, New York, NY, USA. ACM.

Wang, J., Shen, X., Sun, Y., and Qu, A. (2016). Classification with unstructured predictors and an application to sentiment analysis. *Journal of the American Statistical Association*, 111(515):1242–1253.

Wang, T. and Zhao, H. (2017a). A Dirichlet-tree multinomial regression model for associating dietary nutrients with gut microorganisms. *Biometrics*, 73(3):792–801.

Wang, T. and Zhao, H. (2017b). Structured subcomposition selection in regression and its application to microbiome data analysis. *Ann. Appl. Stat.*, 11(2):771–791.

Xia, F., Chen, J., Fung, W. K., and Li, H. (2013). A logistic normal multinomial regression model for microbiome compositional data analysis. *Biometrics*, 69(4):1053–1063.

Yu, G. and Liu, Y. (2016). Sparse regression incorporating graphical structure among predictors. *Journal of the American Statistical Association*, 111(514):707–720.

Zhai, J., Kim, J., Knox, K. S., Twigg, H. L., Zhou, H., and Zhou, J. J. (2018). Variance component selection with applications to microbiome taxonomic data. *Frontiers in Microbiology*, 9:509.

Zhang, T., Shao, M.-F., and Ye, L. (2012). 454 pyrosequencing reveals bacterial diversity of activated sludge from 14 sewage treatment plants. *The ISME Journal*, 6(6):1137–1147.