



Published in final edited form as:

Ann Stat. 2013 June ; 41(3): 1111–1141. doi:10.1214/13-AOS1096.

A LASSO FOR HIERARCHICAL INTERACTIONS

Jacob Bien¹, Jonathan Taylor², and Robert Tibshirani³

Cornell University, Stanford University and Stanford University

Abstract

We add a set of convex constraints to the lasso to produce sparse interaction models that honor the hierarchy restriction that an interaction only be included in a model if one or both variables are marginally important. We give a precise characterization of the effect of this hierarchy constraint, prove that hierarchy holds with probability one and derive an unbiased estimate for the degrees of freedom of our estimator. A bound on this estimate reveals the amount of fitting “saved” by the hierarchy constraint.

We distinguish between *parameter sparsity*—the number of nonzero coefficients—and *practical sparsity*—the number of raw variables one must *measure* to make a new prediction. Hierarchy focuses on the latter, which is more closely tied to important data collection concerns such as cost, time and effort. We develop an algorithm, available in the R package hierNet, and perform an empirical study of our method.

Key words and phrases

Regularized regression; lasso; interactions; hierarchical sparsity; convexity

1. Introduction

There are numerous situations in which additive (main effects) models are insufficient for predicting an outcome of interest. In medical diagnosis, the co-occurrence of two symptoms may lead a doctor to be confident that a patient has a certain disease whereas the presence of either symptom without the other would provide only a moderate indication of that disease. This situation corresponds to a positive (i.e., synergistic) interaction between symptom variables. On the other hand, suppose both symptoms convey redundant information to the doctor about the patient so that knowing both provides no more information about the disease status than either one on its own. This situation is again not additive, but this time

¹Supported in part by the Gerald J. Lieberman Fellowship.

²Supported by NSF Grant DMS-09-06801 and AFOSR Grant 113039.

³Supported by NSF Grant DMS-99-71405 and National Institutes of Health Contract N01-HV-28183.

J. Bien, Department of Biological Statistics and Computational Biology and Department of Statistical Science, Cornell University, Ithaca, New York 14853, USA, jbien@cornell.edu

J. Taylor, Department of Statistics, Stanford University, Stanford, California 94305, USA, jonathan.taylor@stanford.edu

R. Tibshirani, Department of Health, Research, & Policy and Department of Statistics, Stanford University, Stanford, CA 94305, USA, tibs@stanford.edu

SUPPLEMENTARY MATERIAL

Supplement to “A lasso for hierarchical interactions” (DOI: 10.1214/13-AOS1096SUPP;.pdf). We include proofs of Property 1 and of the statement in Remark 3. Additionally, we show that the algorithm for the logistic regression case is nearly identical and give more detail on Algorithm 2.

there is a negative interaction between symptoms. Fitting regression models with interactions is challenging when one has even a moderate number, p , of measured variables, since there are $\binom{p}{k}$ interactions of order k . For this paper, we focus on the case of pairwise ($k = 2$) interaction models, although the ideas we develop generalize naturally to higher-order interaction models.

1.1. Two-way interaction model

We consider a regression model for an outcome variable Y and predictors X_1, \dots, X_p , with pairwise interactions between these predictors. In particular our model has the form

$$Y = \beta_0 + \sum_j \beta_j X_j + \frac{1}{2} \sum_{j \neq k} \Theta_{jk} X_j X_k + \varepsilon, \quad (1)$$

where $\varepsilon \sim N(0, \sigma^2)$. Regardless of whether the predictors are continuous or discrete, we will refer to the additive part as the “main effect” terms and the quadratic part as the “interaction” terms. Our goal is to estimate $\beta \in \mathbb{R}^p$ and $\Theta \in \mathbb{R}^{p \times p}$, where $\Theta = \Theta^T$ and $\Theta_{jj} = 0$. The factor of one half before the interaction summation is a consequence of our notational decision to deal with a symmetric matrix Θ of interactions rather than a vector of length $p(p - 1)/2$. We take $\Theta_{jj} = 0$ throughout this paper because it simplifies notation, but everything carries over if we remove this restriction. Indeed, we provide this as an option in the hierNet (pronounced “hair net”) package.

We observe a training sample, $(x_1, y_1), \dots, (x_n, y_n)$, and our goal is to select a subset of the $p + p(p - 1)/2$ main effect and interaction variables that is predictive of the response, and to estimate the values for the nonzero parameters of the model.

1.2. Strong and weak hierarchy

It is a well-established practice among statisticians fitting (1) to only allow an interaction into the model if the corresponding main effects are also in the model. Such restrictions are known under various names, including “heredity,” “marginality,” and being “hierarchically well-formulated” [Hamada and Wu (1992), Chipman (1996), Nelder (1977), Peixoto (1987)]. There are two types of restrictions, which we will call *strong* and *weak hierarchy*:

$$\begin{aligned} \text{Strong hierarchy: } \hat{\Theta}_{jk} \neq 0 &\Rightarrow \hat{\beta}_j \neq 0 \text{ and } \hat{\beta}_k \neq 0; \\ \text{Weak hierarchy: } \hat{\Theta}_{jk} \neq 0 &\Rightarrow \hat{\beta}_j \neq 0 \text{ or } \hat{\beta}_k \neq 0. \end{aligned}$$

Some statisticians argue that models violating strong hierarchy are not sensible. For example, according to McCullagh and Nelder (1983),

“[T]here is usually no reason to postulate a special position for the origin, so that the linear terms must be included with the cross-term.”

To see that violating strong hierarchy amounts to “postulating a special position for the origin,” consider writing an interaction model as $Y = \beta_0 + (\beta_1 + \Theta_{12}X_2)X_1 + \dots$. First of all,

we would only take $\beta_0 = 0$ if we have special reason to believe that the regression surface must go through the origin. Likewise, taking $\beta_1 = 0$ but $\Theta_{12} \neq 0$ would only be appropriate if we actually believe that X_1 's effect on Y should only be present specifically when X_2 is nonzero. In most situations, we do not think that the variable X_2 that we measured is any more special than $aX_2 + b$. Yet if our model with X_2 violates strong hierarchy, then our model with $aX_2 + b$ (for any $b \neq 0$) is strongly hierarchical. This argument suggests that violations to hierarchy occur in special situations whereas hierarchy is the default.

Another argument in favor of hierarchy has to do with statistical power. In the words of Cox (1984):

“[L]arge component main effects are more likely to lead to appreciable interactions than small components. Also, the interactions corresponding to larger main effects may be in some sense of more practical importance.”

In other words, rather than looking at all possible interactions, it may be useful to focus our search on those interactions that have large main effects. Indeed, the method we propose in this paper makes direct use of this principle.

As a final argument for hierarchy, it is useful to distinguish between two notions of sparsity, which we will call *parameter sparsity* and *practical sparsity*. Parameter sparsity is what most statisticians mean by “sparsity”: the number of nonzero coefficients in the model. Practical sparsity is what someone actually collecting data cares about: the number of variables one needs to *measure* to make predictions at a future time. The hierarchy restriction favors models that “reuse” measured variables whereas a nonhierarchical model does not. The top left panel of Figure 1 gives a small example where this difference is manifest. In fact, a simple calculation shows that this difference can be quite substantial: we can have a hierarchical and a nonhierarchical interaction model with the same parameter sparsity but with the nonhierarchical method having a practical sparsity of $k(k + 1)$ whereas the hierarchical method's practical sparsity is just k .

While taking these arguments to the extreme leads to the use of strong hierarchy exclusively, we develop the case of weak hierarchy in parallel throughout this paper. Weak hierarchy, as the name suggests, can be thought of as a compromise between strong hierarchy and imposing no such structure and appears as a principle in certain statistical methods such as classification and regression trees [Breiman et al. (1984)] and multivariate additive regression splines [Friedman (1991)].

1.3. Sparsity, the lasso and structured sparsity

The lasso [Tibshirani (1996)] is a method that performs both model selection and estimation. It penalizes the squared loss of the data with an ℓ_1 -norm penalty on the parameter vector. This penalty has the property of producing estimates of the parameter vector that are sparse (corresponding to model selection). Given a design matrix $X \in \mathbb{R}^{n \times d}$ and response vector $y \in \mathbb{R}^n$, the lasso is the solution to the convex optimization problem,

$$\text{Minimize}_{\beta_0, \phi} \frac{1}{2} \|y - \beta_0 \mathbf{1} - \tilde{X} \phi\|^2 + \lambda \|\phi\|_1,$$

where $\mathbf{1} \in \mathbb{R}^n$ is the vector of ones. The penalty parameter, $\lambda \geq 0$, controls the relative importance of fitting to the training data (sum-of-squares term) and of sparsity (ℓ_1 penalty term). A natural extension of the lasso to our interaction model (1) would be to take $\phi^T = [\beta^T, \text{vec}(\Theta)^T]$ and $\tilde{X} = (X:Z/2)$, where the columns of $Z \in \mathbb{R}^{n \times p(p-1)}$ correspond to elementwise products of the columns of X . We will refer to this method as the *all-pairs lasso* since it is simply the lasso applied to a data matrix which includes all pairs of interactions (as well as all main effects). It is common with the lasso to standardize the predictors so that they are on the same scale. In this paper, we standardize X so that its columns have mean 0 and standard deviation 1; we then form Z from these standardized predictors and, finally, center the resulting columns of Z . By centering y and \tilde{X} , we may take $\hat{\beta}_0 = 0$.

The lasso's ℓ_1 penalty is neutral to the pattern of sparsity, allowing any sparsity pattern to emerge. The notions of strong and weak hierarchy introduced in Section 1.2 represent situations in which we want to exclude certain sparsity patterns. There has been a growing literature focusing on methods that produce *structured sparsity* [Yuan and Lin (2006), Zhao, Rocha and Yu (2009), Jenatton et al. (2010), Jenatton, Audibert and Bach (2011), Bach (2011), Bach et al. (2012)]. These methods make use of the *group lasso* penalty (and generalizations thereof) which, given a predetermined grouping of the parameters, induces entire groups of parameters to be set to zero [Yuan and Lin (2006)]. Given a set of groups of variables, \mathcal{G} , these methods generalize the ℓ_1 penalty by

$$\sum_{G \in \mathcal{G}} d_G \|\phi_G\|_{\gamma_G},$$

where $\gamma_G > 1$, ϕ_G is ϕ projected onto the coordinates in G , and d_G is a nonnegative weight. Hierarchical structured sparsity is obtained by choosing \mathcal{G} to have nested groups. For example, Zhao, Rocha and Yu (2009) consider the penalty

$$\sum_{j \neq k} \{|\Theta_{jk}| + \|(\beta_j, \beta_k, \Theta_{jk})\|_{\gamma_{jk}}\}.$$

Likewise, the framework of Bach et al. (2012) if specialized to this paper's focus would lead to a penalty of the form

$$\|\Theta\|_1 + \sum_j d_j \|(\Theta_j, \beta_j)\|_q \quad (2)$$

for some $q > 1$ and $d_j > 0$. In fact, Radchenko and James (2010) suggest a penalty for generalized additive models with interactions that reduces to (2) in the linear model case, with $q = 2$ and d_j independent of j .

1.4. This paper

Here, we propose a lasso-like procedure that produces sparse estimates of β and Θ while satisfying the strong or weak hierarchy constraint. In contrast to much of the structured sparsity literature which is based on group lasso penalties, our approach, presented in Section 2, involves adding a set of convex constraints to the lasso. Although we find this form of constraint more naturally interpretable, we show (Remark 3) that this problem can be equivalently expressed in a form that relates it to penalties from the structured sparsity literature such as (2).

A key advantage of our specific choice of penalty structure is that it admits a simple interpretation of the effect of the hierarchy demand. Unlike other hierarchical sparsity methods, which do not pay much attention to the particular choice of norms (as long as $\gamma_G > 1$), our formulation is carefully tailored to allow it to be related directly back to the lasso, permitting one to understand specifically how hierarchy alters the solution (Section 3.1). This feature of our estimator gives it a transparency that exposes the effects (both positive and negative!) of the hierarchy constraint. Furthermore, our characterization suggests that the demand for hierarchy is—analogueous to the demand for sparsity—a form of “regularization.” We develop an unbiased estimator of the degrees of freedom of our method (Section 3.3) and an interpretable upper bound on this quantity, which also points to hierarchy as regularization. In particular, we show that we do not “spend” in degrees of freedom for main effects that are forced into the model by the hierarchy constraint.

Another difference from much of the structured sparsity literature, which aims to develop a broad treatment of structured and hierarchical sparsity methods, is that our focus is narrowed to the problem of interaction models. Our restricted scope allows us to address specifically the performance of such a tool to this important problem. In Section 4, we review previous work on the problem of hierarchical interaction model fitting and selection. These methods fall into three categories: Multi-step procedures, which are defined by an algorithm [Peixoto (1987), Friedman (1991), Turlach (2004), Nardi and Rinaldo (2012), Bickel, Ritov and Tsybakov (2010), Park and Hastie (2008), Wu et al. (2010)]; Bayesian approaches, which specify the hierarchy requirement through a prior [Chipman (1996)]; and, most related to this paper’s proposal, regularized regression methods, which are defined by an optimization problem [Yuan, Joseph and Zou (2009), Zhao, Rocha and Yu (2009), Choi, Li and Zhu (2010), Jenatton et al. (2010), Radchenko and James (2010)]. In Section 5, we study via simulation the statistical implications of imposing hierarchy on an interactions-based estimator under various scenarios (in both the lasso and stepwise frameworks). In Section 6 we present an efficient algorithm for computing our estimator. Real data examples are used to illustrate a distinction we draw between “parameter sparsity” and “practical sparsity” and to discuss hierarchy’s role in promoting the latter.

2. Our proposed method

In Section 1.3, we introduced the *all-pairs lasso*, which can be written as

$$\underset{\beta_0 \in \mathbb{R}, \beta \in \mathbb{R}^p, \Theta \in \mathbb{R}^{p \times p}}{\text{Minimize}} \quad q(\beta_0, \beta, \Theta) + \lambda \|\beta\|_1 + \frac{\lambda}{2} \|\Theta\|_1 \quad \text{s.t.} \quad \Theta = \Theta^T, \quad (3)$$

where $\|\Theta\|_1 = \sum_j \sum_k |\Theta_{jk}|$ and $q(\beta_0, \beta, \Theta)$ is the loss function, typically

$\frac{1}{2} \sum_{i=1}^n (y_i - \beta_0 - x_i^T \beta - \frac{1}{2} x_i^T \Theta x_i)^2 = \frac{1}{2} \|y - \beta_0 \mathbf{1} - X\beta - Z \text{vec}(\Theta)/2\|^2$, but may also include a ridge penalty on the coefficients as discussed later or may be substituted for the binomial negative log-likelihood. The one-half factors in front of terms involving Θ are merely a consequence of the notational choice to represent Θ as a symmetric matrix (with $\Theta_{jj} = 0$ for $j = 1, \dots, p$). In this paper, we propose a modification of the all-pairs lasso that produces models that are guaranteed to be hierarchical.

As motivation for our proposal, consider building hierarchy into the optimization problem as a constraint,

$$\underset{\beta_0 \in \mathbb{R}, \beta \in \mathbb{R}^p, \Theta \in \mathbb{R}^{p \times p}}{\text{Minimize}} \quad q(\beta_0, \beta, \Theta) + \lambda \|\beta\|_1 + \frac{\lambda}{2} \|\Theta\|_1 \quad (4)$$

$$\text{s.t.} \quad \Theta = \Theta^T, \|\Theta_j\|_1 \leq |\beta_j| \quad \text{for } j=1, \dots, p,$$

where Θ_j denotes the j th row (and column, by symmetry) of Θ . Notice that if $\hat{\beta}_j = 0$, then $\|\Theta_j\|_1 > 0$ and $\|\Theta_k\|_1 > 0$ and thus $\hat{\beta}_j = 0$ and $\hat{\beta}_k = 0$. While the added constraints enforce strong hierarchy, they are not convex, which makes (4) undesirable as a method. In this paper, we propose a straightforward convex relaxation of (4), which we call the *strong hierarchical lasso*,

$$\underset{\beta_0 \in \mathbb{R}, \beta^\pm \in \mathbb{R}^p, \Theta \in \mathbb{R}^{p \times p}}{\text{Minimize}} \quad q(\beta_0, \beta^+ - \beta^-, \Theta) + \lambda \mathbf{1}^T (\beta^+ + \beta^-) + \frac{\lambda}{2} \|\Theta\|_1 \quad (5)$$

$$\text{s.t.} \quad \Theta = \Theta^T, \left. \begin{array}{l} \|\Theta_j\|_1 \leq \beta_j^+ + \beta_j^- \\ \beta_j^+ \geq 0, \beta_j^- \geq 0 \end{array} \right\} \text{ for } j=1, \dots, p,$$

where we have replaced the optimization variable $\beta \in \mathbb{R}^p$ by two vectors $\beta^+, \beta^- \in \mathbb{R}^p$. After solving the above problem, our fitted model is of the form $f(\hat{x}) = \hat{\beta}_0 + x^T (\hat{\beta}^+ - \hat{\beta}^-) + x^T \hat{\Theta} \hat{x}/2$. While we might informally think of β^+ and β^- as positive and negative parts of a vector $\beta = \beta^+ - \beta^-$, that is, that $\beta^\pm = \max\{\pm\beta, 0\}$, this is not actually the case since at a solution we can have both $\hat{\beta}_j^+ > 0$ and $\hat{\beta}_j^- > 0$. Indeed, if we were to add the constraints $\beta_j^+ \beta_j^- = 0$ for $j = 1, \dots, p$ to (5), then these would be positive and negative parts and so $\beta_j^+ + \beta_j^- = |\beta_j|$, giving us precisely problem (4). This observation establishes that (5) is a convex relaxation of (4).

The hierarchy constraints can be seen as an embedding into our method of David Cox’s “principle” that “large component main effects are more likely to lead to appreciable interactions than small components.” The constraint

$$\|\Theta_j\|_1 \leq \beta_j^+ + \beta_j^-$$

budgets the total amount of interactions involving variable X_j according to the relative importance of X_j as a main effect. One additional advantage of the convex relaxation is that the constraint is less restrictive. If the best fitting model would have $\|\Theta_j\|_1$ large but $|\beta_j|$ only moderate, this can be accommodated by making β_j^+ and β_j^- both large.

Remark 1—Another possibility for the hierarchy constraint that we have considered is $|\Theta_{jk}| \leq \beta_j^+ + \beta_j^-$; however, we have found that this can lead to an overabundance of interactions relative to main effects.

Remark 2—It is desirable to include in the loss function q an *elastic net* term, $(\varepsilon/2)(\|\Theta\|_F^2 + \|\beta^+\|^2 + \|\beta^-\|^2)$, to ensure uniqueness of the solution [Zou and Hastie (2005)]. We think of $\varepsilon > 0$ as a fixed tiny fraction of λ , such as $\varepsilon = 10^{-8}\lambda$, rather than as an additional tuning parameter. Such a modification does not complicate the algorithm, but simplifies the study of the estimator. In all numerical examples and in the hierNet package, we use this elastic net modification.

Remark 3—We prove in Section 2 of the supplementary materials [Bien, Taylor and Tibshirani (2013)] that (5) may equivalently be written as

$$\begin{aligned} & \text{Minimize} && q(\beta_0, \beta, \Theta) + \lambda \sum_j \max\{|\beta_j|, \|\Theta_j\|_1\} + \frac{\lambda}{2} \|\Theta\|_1. \\ & \beta_0 \in \mathbb{R}, \beta \in \mathbb{R}^p, && \\ & \Theta \in \mathbb{R}^{p \times p}, \Theta = \Theta^T && \end{aligned} \quad (6)$$

This reparameterization of the problem shows its similarities to the group lasso based methods. In place of the more standard penalty $\|(\Theta_j, \beta_j)\|_q$ of (2), we use $\max\{\|\Theta_j\|_1, |\beta_j|\}$. In Section 3.1, we show that this unusual choice of penalty admits a particularly simple interpretation for the effect of imposing hierarchy.

In Section 1.2, we also introduced the notion of weak hierarchy. By simply removing the symmetry constraint on Θ , we get what we call the *weak hierarchical lasso*,

$$\begin{aligned} & \text{Minimize} && q(\beta_0, \beta^+ - \beta^-, \Theta) + \lambda 1^T(\beta^+ + \beta^-) + \frac{\lambda}{2} \|\Theta\|_1 \\ & \beta_0 \in \mathbb{R}, \beta^\pm \in \mathbb{R}^p, \Theta \in \mathbb{R}^{p \times p} && \\ & \text{s.t.} && \left. \begin{aligned} \|\Theta_j\|_1 &\leq \beta_j^+ + \beta_j^- \\ \beta_j^+ &\geq 0, \beta_j^- \geq 0 \end{aligned} \right\} \text{ for } j=1, \dots, p. \end{aligned} \quad (7)$$

Even though at a solution to this problem, $\hat{\Theta}$ is not symmetric, we should ^b think of the interaction coefficient as $(\hat{\Theta}_{jk} + \hat{\Theta}_{kj})/2$ since this is what multiplies the interaction term $x_{ij}x_{ik}$ when computing $f(\hat{x}_i)$.

Remark 4—We can build further on the connection between (2) and (5) discussed in Remark 3. Our removal of the symmetry constraint in (7) is analogous to the technique of duplicating columns of the design matrix used in the overlap group lasso [Obozinski, Jacob and Vert (2011)].

A favorable property that distinguishes our method from previous approaches discussed in Section 4 is the relative transparency of the role that the hierarchy constraint plays in our estimator. This aspect is developed in Section 3.1.

Although our primary focus in this paper is on the Gaussian setting of (1), our proposal extends straightforwardly to other situations, such as the logistic regression setting in which the response is binary. In this case, we simply have $q(\beta_0, \beta, \Theta)$ be the appropriate negative log-likelihood, $-\sum_{i=1}^n y_i \log p_i + (1 - y_i) \log(1 - p_i)$, where

$p_i = [1 + \exp(-\beta_0 - x_i^T \beta - \frac{1}{2} x_i^T \Theta x_i)]^{-1}$. In Section 3 of the supplementary materials [Bien, Taylor and Tibshirani (2013)], we show that solving this problem requires only a minor modification to our primary algorithm. It should also be noted that our estimator (and the algorithms developed to compute it) is designed for both the $p < n$ and $p \geq n$ setting.

As a preliminary example, consider predicting whether a sample of olive oil comes from Southern Apulia based on measurements of the concentration of $p = 8$ fatty acids [Forina et al. (1983)]. The dataset consists of $n = 572$ samples, and we average our results over 100 random equal-sized train-test splits. We compare three methods: (a) a standard lasso with main effects only (MEL), (b) the all-pairs lasso (APL), and (c) the strong hierarchical lasso (HL).

The top left panel of Figure 1 shows an interesting difference between HL and APL. We see that, on average, at a parameter sparsity level of five, the HL model uses four of the measured variables whereas APL uses six. Using the hierarchical model to classify a future olive oil, we only need to measure four rather than six of the fatty acids.

The top right panel of Figure 1 shows the predictive performance (versus the practical sparsity) of the three methods. It appears that HL enjoys the “best of both worlds,” matching the good performance of MEL for low practical sparsity levels (since it tends to pick out the main effects first) and the good performance of APL at high practical sparsity levels (since it can incorporate predictive interactions). Finally, the bottom panel of the figure provides a visual display of a sequence of HL’s solutions (by varying λ). Nonzero main effects are shown as filled nodes, and edges indicate nonzero interactions. Since all edges are incident to filled nodes, we see that strong hierarchy holds.

In the next section, we present several properties of our estimator that shed light on the effect of adding the convex hierarchy constraint to the lasso. Among these properties is an unbiased estimate of the degrees of freedom of our estimator. We view this degrees-of-freedom result as valuable primarily for the sake of understanding the effect of hierarchy. While such an estimate could be used for parameter selection, we prefer cross validation to select λ since this is more directly tied to the goal of prediction.

3. Properties

3.1. Effect of the constraint

A key advantage of formulating an estimator as a solution to a convex problem is that it can be completely characterized by a set of optimality conditions, known as the Karush–Kuhn–

Tucker (KKT) conditions. These conditions are useful for understanding the effect that the hierarchy constraint in (5) and (7) has on our solutions. In this section, we will study the simplest case, taking $q(\beta_0, \beta, \Theta)$ to be the quadratic loss function with no elastic net penalty. We let

$$\begin{aligned} r^{(-j)} &= y - \hat{y} + x_j \hat{\beta}_j, \\ r^{(-jk)} &= y - \hat{y} + (x_j * x_k)(\hat{\Theta}_{jk} + \hat{\Theta}_{kj})/2 \end{aligned}$$

denote partial residuals (where $*$ denotes elementwise multiplication, \hat{y} the vector of fitted values and x_j the j th predictor), and we assume that $\|x_j\|_2 = 1$. For linear regression, the KKT conditions are known as the normal equations and can be written as

$$\hat{\beta}_j = x_j^T r^{(-j)}, \quad \hat{\Theta}_{jk} = \frac{(x_j * x_k)^T r^{(-jk)}}{\|x_j * x_k\|^2}.$$

The all-pairs lasso solution satisfies

$$\hat{\beta}_j = \mathcal{S}(x_j^T r^{(-j)}, \lambda), \quad \hat{\Theta}_{jk} = \frac{\mathcal{S}[(x_j * x_k)^T r^{(-jk)}, \lambda]}{\|x_j * x_k\|^2}, \quad (8)$$

where \mathcal{S} denotes the soft-thresholding operator defined by $\mathcal{S}(c, \lambda) = \text{sign}(c)(|c| - \lambda)_+$. Written this way, we see that the lasso is similar to linear regression, but all coefficients are shrunken toward 0, with some coefficients (those for which $|x_j^T r^{(-j)}| \leq \lambda$) set to zero. It is instructive to examine the corresponding statements for the strong and weak hierarchical lasso methods.

Property 1—The coefficients of the strong and weak hierarchical lassos with $\lambda > 0$ and taking $q(\beta_0, \beta, \Theta)$ to be the quadratic loss (with no elastic net penalty) satisfy:

- Strong:

$$\begin{aligned} \hat{\beta}_j^+ - \hat{\beta}_j^- &= \mathcal{S}(x_j^T r^{(-j)}, \lambda - \hat{\alpha}_j), \\ \hat{\Theta}_{jk} &= \frac{\mathcal{S}[(x_j * x_k)^T r^{(-jk)}, \lambda + \hat{\alpha}_j + \hat{\alpha}_k]}{\|x_j * x_k\|^2}; \end{aligned}$$

- Weak:

$$\begin{aligned} \hat{\beta}_j^+ - \hat{\beta}_j^- &= \mathcal{S}(x_j^T r^{(-j)}, \lambda - \tilde{\alpha}_j), \\ \frac{\hat{\Theta}_{jk} + \hat{\Theta}_{kj}}{2} &= \frac{\mathcal{S}[(x_j * x_k)^T r^{(-jk)}, \lambda + 2 \min\{\tilde{\alpha}_j, \tilde{\alpha}_k\}]}{\|x_j * x_k\|^2} \end{aligned}$$

for some $\hat{\alpha}_j \geq 0, j = 1, \dots, p$ with $\hat{\alpha}_j = 0$ when $\|\hat{\Theta}_j\|_1 < \hat{\beta}_j^+ + \hat{\beta}_j^-$ (and likewise for $\tilde{\alpha}_j$).

Proof: See Section 1 of the supplementary materials [Bien, Taylor and Tibshirani (2013)].

The $\hat{\alpha}_j, \tilde{\alpha}_j$ appearing in the above two properties are optimal dual variables corresponding to the j th hierarchy constraint for the strong and weak hierarchical lasso problems, respectively. When $\|\hat{\Theta}_j\|_1 < \hat{\beta}_j^+ + \hat{\beta}_j^-$, we have $\hat{\alpha}_j = 0$ (or $\tilde{\alpha}_j = 0$) by complementary slackness. Comparing these expressions to those of the all-pairs lasso gives insight into the effect of the constraint. Property 1 reveals that the overall form of the all-pairs lasso and hierarchical lasso methods is identical. The difference is that the hierarchy constraint leads to a reduction in the shrinkage of certain main effects and an increase in the shrinkage of certain interactions. In particular, we see that when the hierarchy constraints are loose at the solution, that is, $\|\hat{\Theta}_j\|_1 < \hat{\beta}_j^+ + \hat{\beta}_j^-$, the weak hierarchical lasso's optimality conditions become identical to the all-pairs lasso (since $\hat{\alpha}_j = 0$) for all coefficients involving x_j . For the strong hierarchical lasso, when both the j th and k th constraints are loose, the optimality conditions match those of the all-pairs lasso for the coefficients of x_j, x_k and $x_j * x_k$. The methods differ when constraints are active, that is, when $\|\hat{\Theta}_j\|_1 = \hat{\beta}_j^+ + \hat{\beta}_j^-$, which allows $\hat{\alpha}_j$ (or $\tilde{\alpha}_j$) to be nonzero. Intuitively, this case corresponds to the situation in which hierarchy would not have held "naturally" (i.e., without the constraint), and the corresponding dual variable plays the role of reducing $\hat{\Theta}_j$ in ℓ_1 -norm and increasing $\hat{\beta}_j^+ + \hat{\beta}_j^-$ until the constraint is satisfied. The way in which the weak and strong hierarchical lasso methods perform this shrinkage is different, but both are identical to the all-pairs lasso when all constraints are loose.

3.2. Hierarchy guarantee

In Section 2, we showed that adding the constraint $\|\Theta_j\|_1 = |\beta_j|$ would guarantee that hierarchy holds. However, we have not yet shown that the same is true of the convex relaxation's constraint, $\|\Theta_j\|_1 \leq \beta_j^+ + \beta_j^-$. In particular, while $\hat{\Theta}_{jk} \neq 0 \Rightarrow \hat{\beta}_j^+ + \hat{\beta}_j^- \neq 0$, we could still have $\hat{\beta}_j^+ - \hat{\beta}_j^- = 0$. This would correspond to a model in which $X_j X_k$ is used in the model, but X_j is not. Intuitively, we would expect that if $\hat{\beta}_j^+ > 0$, then $\hat{\beta}_j^+ = \hat{\beta}_j^-$ is analogous to getting an exact zero in linear regression (i.e., a zero probability event). In this section, we establish that this is in fact the case.

In particular, we study (5) and (7) where $q(\beta_0, \beta, \Theta)$ includes an elastic net term. The importance of this modification is that it ensures uniqueness, simplifying the analysis. As noted in Remark 2, we think of ε as a small, fixed proportion of λ rather than as a separate tuning parameter.

Property 2—Suppose y is absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^n . If $(\hat{\beta}^+, \hat{\beta}^-, \hat{\Theta})$ solves (5), where $q(\beta_0, \beta, \Theta)$ is the quadratic loss with an $\varepsilon > 0$ ridge penalty, then strong hierarchy holds with probability 1, that is,

$$\hat{\Theta}_{jk} \neq 0 \Rightarrow \hat{\beta}_j^+ - \hat{\beta}_j^- \neq 0 \quad \text{and} \quad \hat{\beta}_k^+ - \hat{\beta}_k^- \neq 0.$$

Proof: See Appendix A.

To understand how dropping the symmetry constraint leads to the “or” statement required of weak hierarchy, note that $X_j X_k$ is in the weak hierarchical lasso model if and only if $\hat{\Theta}_{jk} + \hat{\Theta}_{kj} = 0$. This holds only if $\hat{\Theta}_{jk} = 0$ or $\hat{\Theta}_{kj} = 0$.

Property 3—Suppose y is absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^n . If $(\hat{\beta}^+, \hat{\beta}^-, \hat{\Theta})$ solves (7), where $q(\beta_0, \beta, \Theta)$ is the quadratic loss with an $\varepsilon > 0$ ridge penalty, then weak hierarchy holds with probability 1, that is,

$$\frac{\hat{\Theta}_{jk} + \hat{\Theta}_{kj}}{2} \neq 0 \Rightarrow \hat{\beta}_j^+ - \hat{\beta}_j^- \neq 0 \quad \text{or} \quad \hat{\beta}_k^+ - \hat{\beta}_k^- \neq 0.$$

Proof: See Appendix A.

3.3. Degrees of freedom

In classical statistics, the degrees of freedom of a procedure refer to the dimension of the space over which its fitted values can vary. It is useful in that it provides a measure of how much “fitting” the procedure is doing. This notion can be generalized to adaptive procedures such as the lasso [Stein (1981), Efron (1986), Efron et al. (2004), Zou, Hastie and Tibshirani (2007)]. See (Ryan Tibshirani and Taylor (2012) for a thorough discussion. If given data $y \in \mathbb{R}^n$, a procedure h produces fitted values $\hat{y} = h(y) \in \mathbb{R}^n$, the degrees of freedom of the procedure h is defined to be

$$df(h) = \frac{1}{\sigma^2} \sum_{i=1}^n \text{cov}(y_i, \hat{y}_i). \quad (9)$$

Property 4—Suppose $y \sim N(\mu, \sigma^2 I_n)$. An unbiased estimate of the degrees of freedom of the strong hierarchical lasso, with quadratic loss and no ridge penalty, is given by

$$\hat{df}_\lambda = \text{rank}(\tilde{X}P),$$

where $\tilde{X} = (X: -X:Z/2: -Z/2)$ with Z containing the interactions, and P is a projection matrix which depends on the sign pattern of $(\hat{\beta}^+, \hat{\beta}^-, \hat{\Theta})$ and on the set of hierarchy constraints that are tight.

Proof: See Appendix B.

Figure 2 provides a numerical evaluation of how well \hat{df}_λ estimates df_λ . We fix $X \in \mathbb{R}^{n \times p}$, $\beta \in \mathbb{R}^p$ and $\Theta \in \mathbb{R}^{p \times p}$, and we generate $B = 10,000$ Monte Carlo replicates $y^{(1)}, \dots, y^{(B)} \in \mathbb{R}^n$. For each replicate, we fit the strong hierarchical lasso along a grid of λ values to get

$(\hat{\beta}_\lambda^{+(b)}, \hat{\beta}_\lambda^{-(b)}, \hat{\Theta}_\lambda^{(b)})$ and $\hat{y}_\lambda^{(b)} \in \mathbb{R}^n$. From these values, we compute Monte Carlo estimates of df_λ from the definition in (9) and of $E[\hat{df}_\lambda]$.

While \hat{df}_λ can be calculated from the data and is therefore useful as an unbiased way of calibrating the amount of fitting the strong hierarchical lasso is doing, this expression is difficult to interpret. However, it turns out that we can *bound* \hat{df}_λ by a quantity that does make more sense:

Property 5—Let $\mathcal{T} = \{j: \|\hat{\Theta}_j\|_1 = \hat{\beta}_j^+ + \hat{\beta}_j^-\}$, $\mathcal{A}_\beta = \{j: \hat{\beta}_j^+ - \hat{\beta}_j^- \neq 0\}$, $\mathcal{A}_\pm = \{j: \hat{\beta}_j^\pm > 0\}$ and $\mathcal{A}_\Theta = \{jk: \Theta_{jk} > 0, j < k\}$. Then,

$$\hat{df}_\lambda \leq |\mathcal{A}_\beta| + |\mathcal{A}_\Theta| - |\mathcal{T} \cap (\mathcal{A}_+ \Delta \mathcal{A}_-)|$$

holds almost surely, where $\mathcal{A}_+ \Delta \mathcal{A}_- = (\mathcal{A}_+ \setminus \mathcal{A}_-) \cup (\mathcal{A}_- \setminus \mathcal{A}_+)$.

Proof: See Appendix B.

By contrast, for the all-pairs lasso in the case that $p + \binom{p}{2} < n$ and the design matrix is full rank, we have $df_\lambda(\text{APL}) = E[|\mathcal{A}_\beta| + |\mathcal{A}_\Theta|]$ [Zou, Hastie and Tibshirani (2007)]. In other words, the strong hierarchical lasso does not “pay” (in terms of fitting) for those main effects, $\hat{\beta}_j^+ - \hat{\beta}_j^-$, that are forced into the model by the hierarchy constraint to accommodate a strong interaction. Notice that we do pay for a nonzero main effect if *both* $\hat{\beta}_j^+$ and $\hat{\beta}_j^-$ are nonzero. This makes sense since the constraint could be satisfied with just one of these variables nonzero, but in this case it is advantageous to the fit to make both nonzero. In Figure 2, we find that this bound is in expectation visually indistinguishable from $E[\hat{df}_\lambda]$.

4. Related work

There has been considerable interest in fitting interaction models in statistics and related fields. We focus here on an overview of methods that aim at forming predictive models that satisfy the hierarchical interactions restriction.

4.1. Multi-step procedures

Many statistics textbooks discuss a simple stepwise procedure in which one iteratively considers adding or removing the “best” variable (whether it be main effect or interaction); they add that one should only consider including an interaction if its main effects are in the model [e.g., see backward elimination in Agresti (2002), Section 6.1.3]. In doing so, they are enforcing the strong hierarchy restriction. Such procedures are ubiquitous [Nelder (1997), Peixoto (1987)] as are more recent versions [Friedman (1991), Bickel, Ritov and Tsybakov (2010), Park and Hastie (2008), Wu et al. (2010)]. Another approach is to perform model selection first without considering hierarchy and then to include any lower-order terms

necessary to satisfy hierarchy as a post-processing step [Nardi and Rinaldo (2012)]. Finally, Turlach (2004) and Yuan, Joseph and Lin (2007) consider modifying the LARS algorithm [Efron et al. (2004)] so that hierarchy is enforced.

4.2. Bayesian approaches

Another set of procedures for building hierarchical interaction models comes from a Bayesian viewpoint. Chipman (1996) adapts the *stochastic search variable selection* (SSVS) approach of George and McCulloch (1993) to produce strong or weak hierarchical interaction models. SSVS makes use of a hierarchical normal mixture model to perform variable selection in regression. Every variable has a latent binary variable indicating whether it is “active.” Conditional on this latent variable, each coefficient is a 0-mean normal with variance determined by the latent importance of the coefficient. The original SSVS paper chooses a prior in which the importance of each variable is an independent Bernoulli. Chipman (1996) introduces dependence into the prior so that Θ_{jk} is important only if β_j and/or β_k is important as well.

4.3. Optimization-based approaches

Choi, Li and Zhu (2010) formulate a nonconvex optimization problem to get sparse hierarchical interaction models. They write $\Theta_{jk} = \Gamma_{jk}\beta_j\beta_k$, where β are the main effect coefficients and then apply ℓ_1 penalties on β and Γ . Notice that $\Theta_{jk} = 0$ implies $\beta_j = 0$ and $\beta_k = 0$. The nonconvexity arises in writing Θ_{jk} as the product of optimization variables.

Most similar to this paper’s proposal is a series of methods which formulate convex optimization problems to give sparse hierarchical interaction models. Yuan, Joseph and Zou (2009) modify the nonnegative garrote [Breiman (1995)] by adding linear inequality constraints to enforce hierarchy. In this sense, our method can be seen as the adaption of their approach to the lasso.

Finally, as discussed in Section 1.3, another set of convex methods makes use of the group lasso penalty [Yuan and Lin (2006)]. Zhao, Rocha and Yu (2009) [and, relatedly, Jenatton, Audibert and Bach (2011)] describe composite absolute penalties (CAP), a very broad class of penalties that can achieve group and hierarchical sparsity. To achieve “hierarchical selection,” they put forward the principle that a penalty of the form $\|(\varphi_1, \varphi_2)\|_\gamma + |\varphi_1|$, with $\gamma > 1$, induces φ_2 to be zero only when φ_1 is zero as well. For hierarchical interaction models, they suggest a penalty of the form $\lambda \sum_{j < k} [|\Theta_{jk}| + \|(\beta_j, \beta_k, \Theta_{jk})\|_{\gamma, k}]$. This framework has been developed in the structured sparsity literature [e.g., Bach et al. (2012)]. Radchenko and James (2010) introduce VANISH, which uses this nested-group principle to achieve hierarchical sparsity in the context of nonlinear interactions. Their penalty in the setting of (1) is $\sum_j [\lambda_1 \|(\beta_j, \Theta_j)\|_2 + \lambda_2 \|\Theta_j\|_1]$. As noted in Remark 3, our proposal is closer to CAP and VANISH than it may first appear. Our problem can be rewritten to have a penalty of the form $\lambda \sum_j [\max\{|\beta_j|, \|\Theta_j\|_1\} + (1/2)\|\Theta_j\|_1]$. In this sense, the penalty is in the spirit of CAP and related methods although it does not quite fall into the class of CAP (since ours involves a sum of norms of norms). It is most similar to VANISH in that it combines all of Θ_j into the term involving β_j .

5. Empirical study

5.1. Simulations

Our main interest in this section is to study the advantages and disadvantages of restricting one’s interaction models to those that honor hierarchy. Clearly, the effectiveness of such a strategy depends on the true model generating the data. We take $n = 100$ and $p = 30$ (435 two-way interactions) and consider four scenarios:

- I.** Truth is hierarchical: $\Theta_{jk} = 0 \Rightarrow \beta_j = 0, \beta_k = 0$;
- II.** Truth is anti-hierarchical: $\Theta_{jk} = 0 \Rightarrow \beta_j = 0, \beta_k = 0$;
- III.** Truth only has interactions: $\beta_j = 0$ for all j ;
- IV.** Truth only has main effects: $\Theta_{jk} = 0$ for all jk .

In cases (I), (II), (IV), we set 10 elements of β to be nonzero (with random sign), and, in cases (I), (II), (III), we set 20 elements of the submatrix of $\Theta = \Theta^T$ to be nonzero. The signal-to-noise ratio (SNR) for the main effects part of the signal is about 1.5 whereas the SNR for the interactions part is about 1.

We study the effectiveness of the hierarchy constraint in the context of both the lasso and forward stepwise regression. Forward stepwise regression refers to a greedy strategy for generating a sequence of linear regression models in which we start with an intercept-only model and then at each step add the variable that leads to the greatest decrease in the residual sum of squares. We choose forward stepwise as a basis of comparison since it has a simple modification that we think may be the hierarchical interactions approach most commonly used by statisticians. The modification is to restrict the set of interactions that could be added at a given step to only those between main effect variables currently in the model. A backward stepwise version of this approach is suggested in Peixoto (1987).

We compare six methods, corresponding to each cell of the following table:

	Hierarchical	All-pairs	Main effects only
Lasso	HL (our method)	APL	MEL
Fwd stepwise	HF	APF	MEF

Each method has a single tuning parameter: for the lasso methods, the penalty parameter, λ , and for the forward stepwise methods, the number of variables, k . We fit each method along a grid of tuning parameter values and select the model with the smallest mean squared error, $E\|\hat{y} - \mu\|^6$. Note that such an operation is only possible in simulation since it requires knowing μ ; however, doing so avoids the added variance of cross validation without being biased in favor of any particular method. The results presented are based on 100 simulations from the underlying model. Figure 3 shows the expected prediction error, $\sigma^6 + E[(\hat{y} - \mu)^6]$. Panel (I) shows that when the truth is hierarchical, methods that assume hierarchy (HL, HF) do better than the rest. These methods have “concentrated” their power on the correct set of models and therefore receive the biggest payoff for being correct. APL does better than

MEL and MEF since it succeeds in incorporating some of the correct interactions (recall that interactions make up one quarter of the signal). In panel (II), we notice our first surprise—that HL predicts well relative to the others even when the truth is not hierarchical! We would have expected APL (or APF) to be the clear winner in this situation since surely the hierarchy assumption can only be detrimental in this “anti-hierarchical” scenario. The reason APL does not outperform HL in this scenario is because APL has trouble identifying the main effects (it gets swamped by the 435 interaction variables). In light of Section 3.1, this is where the hierarchy constraint helps—main effects are penalized by less and interactions by more. Even though APL is better able to find the correct interactions than HL, as seen in panel (II) of Figure 4, APL does not predict as well as HL because it fails to find the main effects, which constitute three quarters of the signal. Relatedly, in a “hierarchical truth” scenario similar to (I) but with $p > n$ (not presented here), we have in fact observed MEL doing better than APL (though not as well as HL) since APL is not able to detect interactions accurately enough to make up for its inferior ability to detect main effects. By contrast, HL does best in that scenario, aided by hierarchy to capture both the main effect and interaction components of the signal.

In panel (III), we see a situation where APL does dominate HL. Since there are no main effects in the signal, all that is relevant is a method’s ability to find the interactions. HL identifies fewer correct interactions than APL since any main effect “information” that HL is using is spurious. Finally in panel (IV), we see a situation where MEF, HF, MEL do better than the rest. Here again we find that the hierarchy methods beat the all-pairs methods since they favor main effects.

It is particularly illuminating to note the difference in performance between HL and HF. HF in scenarios (II), (III) and (IV) performs very similarly to the main effect only models. In (II) and (III), HL does much better than HF both in terms of prediction error and in ability to correctly identify interactions. HL appears to be far less sensitive to violations of hierarchy than HF. This difference is attributable to the joint nature in which HL acts: the decision to include a main effect is made at the same time as decisions about interactions. This allows a strong interaction to “pull” itself into the model. By contrast, HF selects main effects with no regard to the information contained in the interactions.

5.2. Data examples

Rhee et al. (2006) study six nucleoside reverse tran-scriptase inhibitors (NRTIs) that are used to treat HIV-1. The target of these drugs can become resistant through mutation, and Rhee et al. (2006) compare a collection of models for predicting these drug’s (log) susceptibility—a measure of drug resistance—based on the location of mutations. In the six cases, there are between $p = 211$ and $p = 218$ sites with mutations occurring in the $n = 784$ to $n = 1073$ samples. While they focus on main effect only models, we consider here the all-pairs lasso (APL) and weak hierarchical lasso (HL) in addition to the standard main effects lasso (MEL). We train on half of the samples and test on the remaining samples. To reduce the dependence of the results on the particular random training-test split, we repeat this process twenty times and average the results. Figure 5 shows the average test RMSE versus the average practical sparsity for each of the six drugs. In all cases but ABC, we find that

HL achieves a better test error at most levels of practical sparsity than APL. That said, if the number of mutations one has to measure is not of concern (so that we can choose for each method the minimum RMSE model), then no method dominates in all the situations. It is worth conceding—since this is a paper on interactions—that in several of the cases a pure main effects model appears to be the best option.

6. Algorithmics

Some of the fastest lasso solvers rely on coordinate descent, which amounts to iteratively applying (8) until convergence [Friedman, Hastie and Tibshirani (2010)]. Tseng (2001) proves that blockwise coordinate descent converges to the global minimum for a convex problem specifically when the nondifferentiable part of the problem is blockwise separable. In the case of the strong hierarchical lasso, the hierarchy constraints combined with the symmetry constraint couple all the parameters together, meaning that coordinate descent is prone to getting stuck at sub-optimal points. To see this, note that $\Theta_{jk} = \Theta_{kj}$ appears in two constraints $\|\Theta_j\|_1 \leq \beta_j^+ + \beta_j^-$ and $\|\Theta_k\|_1 \leq \beta_k^+ + \beta_k^-$. By contrast, the constraints in the weak hierarchical lasso problem are blockwise separable so that blockwise coordinate descent on blocks of the form $(\Theta_j, \beta_j^+, \beta_j^-)$ for $j = 1, \dots, p$ does work. We begin by discussing our approach to solving the weak hierarchical lasso problem. In Section 6.2 we discuss how we can solve a sequence of weak hierarchical lasso problems that converges to a solution of the strong hierarchical lasso.

6.1. Solving the weak hierarchical lasso

While blockwise coordinate descent would work for solving the weak hierarchical lasso problem, we instead describe a generalized gradient descent approach. Given a problem of the form

$$\text{Minimize}_{\phi} g(\phi) + h(\phi), \quad (10)$$

in which g is convex and differentiable with a Lipschitz gradient and h is convex, generalized gradient descent works by solving a sequence of problems of the form

$$\hat{\phi}_k \leftarrow \arg \min_{\phi} \frac{1}{2t} \|\phi - [\hat{\phi}_{k-1} - t \nabla g(\hat{\phi}_{k-1})]\|^2 + h(\phi),$$

where t is a suitably chosen step size [Beck and Teboulle (2009)]. These subproblems are easier to solve than (10) since they replace g by a spherical quadratic. Under the previously stated conditions, generalized gradient descent is guaranteed to get within $O(1/k)$ of the optimal value after k steps; in fact, with a simple modification to the algorithm, this rate improves to $O(1/k^6)$ [Beck and Teboulle (2009)]. Looking back at (7), we take g to be the differentiable part, $q(\beta_0, \beta^+ - \beta^-, \Theta) + \lambda 1^T (\beta^+ + \beta^-)$ and h to be the ℓ_1 penalty on Θ and the set of constraints. The subproblem is of the form

$$\begin{aligned} & \text{Minimize}_{\beta^\pm \in \mathbb{R}^p, \Theta \in \mathbb{R}^{p \times p}} \frac{1}{2t} \|\beta^+ - \tilde{\beta}^+\|^2 + \frac{1}{2t} \|\beta^- - \tilde{\beta}^-\|^2 + \frac{1}{2t} \|\Theta - \tilde{\Theta}\|_F^2 + \frac{\lambda}{2} \|\Theta\|_1 \\ & \text{s.t.} \quad \left. \begin{aligned} \|\Theta_j\|_1 &\leq \beta_j^+ + \beta_j^- \\ \beta_j^+ &\geq 0, \beta_j^- \geq 0 \end{aligned} \right\} \text{ for } j=1, \dots, p, \end{aligned}$$

where $(\beta^+, \beta^-, \Theta)$ depends on the previous iteration's solution and the data, X , Z and y . The exact form of $(\tilde{\beta}^+, \tilde{\beta}^-, \tilde{\Theta})$ is given in Algorithm for solving (7), where $q(\beta_0, \beta, \Theta)$ includes an elastic net penalty as described in Remark 2 of Section 2. The above problem decouples into p separate pieces involving $(\Theta_j, \beta_j^+, \beta_j^-)$ that could be solved in parallel:

$$\begin{aligned} & \text{Minimize}_{\beta_j^\pm \in \mathbb{R}, \Theta_j \in \mathbb{R}^{p-1}} \frac{1}{2t} (\beta_j^+ - \tilde{\beta}_j^+)^2 + \frac{1}{2t} (\beta_j^- - \tilde{\beta}_j^-)^2 + \frac{1}{2t} \|\Theta_j - \tilde{\Theta}_j\|^2 + \frac{\lambda}{2} \|\Theta_j\|_1 \\ & \text{s.t.} \quad \|\Theta_j\|_1 \leq \beta_j^+ + \beta_j^-, \beta_j^+ \geq 0, \beta_j^- \geq 0. \end{aligned} \tag{11}$$

Algorithm 1

WEAK-HIERNET: Generalized gradient descent to solve weak hierarchical lasso, (7), with elastic net penalty ε .

Inputs: $X \in \mathbb{R}^{n \times p}$, $Z \in \mathbb{R}^{n \times (p-1)}$, $\lambda > 0$. Initialize $(\beta^{(0)}, \tilde{\beta}^{(0)}, \Theta^{(0)})$.

For $k = 1, 2, \dots$ until convergence:

 Compute residual: $r^{(k-1)} \leftarrow y - X(\beta^{\hat{(k-1)}} - \tilde{\beta}^{(k-1)}) - Z\Theta^{\hat{(k-1)}}$.

 For $j = 1, \dots, p$:

$$\begin{aligned} (\hat{\beta}_j^{+(k)}, \hat{\beta}_j^{-(k)}, \hat{\Theta}_j^{(k)}) &\leftarrow \text{ONEROW}(\delta \hat{\beta}_j^{+(k-1)} - tX_j^T r^{(k-1)}, \\ &\quad \delta \hat{\beta}_j^{-(k-1)} + tX_j^T r^{(k-1)}, \\ &\quad \delta \hat{\Theta}_j^{(k-1)} - tZ_{(j,\cdot)}^T \hat{\Theta}_j^{(k-1)}), \end{aligned}$$

where ONEROW is given in Algorithm 3, $\delta = 1 - t\varepsilon$, and $Z_{(j,\cdot)} \in \mathbb{R}^{n \times (p-1)}$ denotes the columns of Z involving X_j .

In Appendix C, we derive an algorithm, ONEROW, that solves (11) based on the observation that, in terms of an optimal dual variable α , a solution is simply $\hat{\Theta}_j = \mathcal{S}[\Theta_j, t(\lambda/2 + \alpha)]$ and $\hat{\beta}_j^\pm = [\hat{\beta}_j^\pm + t\hat{\alpha}]_+$.

We solve (7) along a sequence of λ values, from large to small, using the solution from the previous λ as a warm start for the next. The WEAK-HIERNET algorithm gets within ε of the optimal value of (7) in $O(p^6 \max\{n, p\}/\varepsilon)$ time.

6.2. Solving the strong hierarchical lasso

In Section 6.1, we noted that each step of generalized gradient descent conveniently decouples into p single-variable optimization problems. However, for the strong hierarchical lasso, (5), the symmetry constraint ties all variables together. We therefore make use of Alternating Direction Method of Multipliers (ADMM), which is a very widely applicable

framework that allows convex problems to be split apart into separate easier subproblems [Boyd et al. (2011)].

Given a convex problem of the form Minimize $_{\varphi} f(\varphi) + g(\varphi)$, we rewrite it equivalently as Minimize $_{\varphi} \varphi f(\varphi) + g(\varphi)$ s.t. $\varphi = \varphi$, and then the ADMM algorithm repeats the following three steps until convergence:

1. $\hat{\varphi} \leftarrow \operatorname{argmin}_{\varphi} [f(\varphi) + (\rho/2)\|\varphi - \hat{\varphi} + \hat{u}/\rho\|^6]$.
2. $\hat{\varphi} \leftarrow \operatorname{argmin}_{\varphi} [g(\varphi) + (\rho/2)\|\varphi - \hat{\varphi} - \hat{u}/\rho\|^6]$.
3. $\hat{u} \leftarrow \hat{u} + \rho(\hat{\varphi} - \hat{\varphi})$.

Thus the ADMM algorithm separates the two difficult parts of the problem, f and g , into separate optimization problems. The dual variable \hat{u} serves to pull these two problems together, resulting in an algorithm that is guaranteed to converge to a solution as long as $\rho > 0$. In practice, the value of ρ affects the speed of convergence.

In our case, we use ADMM to separate the hierarchy constraints, involving $(\beta^+, \beta^-, \Theta)$ from the symmetry constraint, which will involve a symmetric version of Θ , which we call Ω :

$$\begin{aligned} & \text{Minimize}_{\beta_0 \in \mathbb{R}, \beta^{\pm} \in \mathbb{R}^p, \Theta, \Omega \in \mathbb{R}^{p \times p}} q(\beta_0, \beta^+ - \beta^-, \Theta) + \lambda 1^T (\beta^+ + \beta^-) + \frac{\lambda}{2} \|\Theta\|_1 \\ & \text{s.t. } \left. \begin{array}{l} \Omega = \Omega^T, \Theta = \Omega \\ \|\Theta_j\|_1 \leq \beta_j^+ + \beta_j^- \\ \beta_j^+ \geq 0, \beta_j^- \geq 0 \end{array} \right\} \text{ for } j=1, \dots, p. \end{aligned} \tag{12}$$

The resulting ADMM algorithm is given in Algorithm 2, which is explained in greater detail in Section 4 of the supplementary materials [Bien, Taylor and Tibshirani (2013)]. Conceptually, the algorithm alternately updates two matrices, Θ and Ω . Throughout the algorithm, we update Θ by solving a version of problem (7), and we update Ω by symmetrizing a version of Θ . At convergence, $\hat{\Theta} = \hat{\Omega}$, and thus $\hat{\Theta}$ is both symmetric and satisfies the hierarchy constraints.

Algorithm 2

STRONG-HIERNET: Solve (5) via ADMM.

Inputs: $X \in \mathbb{R}^{n \times p}, Z \in \mathbb{R}^{n \times p(p-1)}, \lambda > 0, \rho > 0$.

Initialize $(\hat{\beta}^+, \hat{\beta}^-, \hat{\Theta}), \hat{\Omega}, \hat{U}$.

Repeat until convergence:

- 1 WEAK-HIERNET(X, Z, λ), but in the call to ONEROW replace the argument $\delta \hat{\Theta}_j^{(k-1)} - t Z_{(j,\cdot)}^T \hat{r}^{(k-1)}$ with $\delta \hat{\Theta}_j^{(k-1)} - t Z_{(j,\cdot)}^T \hat{r}^{(k-1)} + \rho [\hat{\Theta}_j^{(k-1)} - \hat{\Omega}_j] + \hat{U}_j$. Also, initialize with $(\hat{\beta}^+, \hat{\beta}^-, \hat{\Theta})$.
 - 2 $\hat{\Omega} \leftarrow \frac{1}{2}(\hat{\Theta} + \hat{\Theta}^T) + \frac{1}{2\rho}(\hat{U} + \hat{U}^T)$.
 - 3 $\hat{U} \leftarrow \hat{U} + \rho(\hat{\Theta} - \hat{\Omega})$
-

7. Discussion

In this paper, we have proposed a modification to the lasso for fitting strong and weak hierarchical interaction models. These two approaches are closely tied, and our algorithms to solve the two exploit their similar structure. A key advantage of our framework is that it admits a simple characterization of the effect of imposing hierarchy. We compare our hierarchical methods to the lasso and to stepwise procedures to understand the implications of demanding hierarchy. We introduce a distinction between models that have a small number of parameters and those that require measuring only a small number of variables. The hierarchical interaction requirement favors models with the latter type of sparsity, a feature that is desirable when performing measurements is costly, time consuming, or otherwise inconvenient. The R package *hierNet* provides implementations of our strong and weak methods, both for Gaussian and logistic losses. This work has potential applications to genomewide association studies. In future work, we intend to extend this framework to contexts in which only certain interactions should be considered such as in gene-environment interaction models.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Will Fithian, Max Graziop-Sell, Brad Klingenberg, Balasubramanian Narasimhan and Ryan Tibshirani for useful conversations and two referees and an Associate Editor for helpful comments.

References

- Agresti, A. *Categorical Data Analysis*. 2. Wiley-Interscience; New York: 2002.
- Bach F. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*. 2011; 4:1–106.
- Bach F, Jenatton R, Mairal J, Obozinski G. Structured sparsity through convex optimization. *Statist Sci*. 2012; 27:450–468.
- Beck A, Teboulle M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J Imaging Sci*. 2009; 2:183–202.
- Bickel, P.; Ritov, Y.; Tsybakov, A. *Inst Math Stat Collect*. Vol. 6. Inst. Math. Statist; Beachwood, OH: 2010. Hierarchical selection of variables in sparse high-dimensional regression. *Borrowing Strength: Theory Powering Applications—A Festschrift for Lawrence D. Brown*; p. 56-69.
- Bien J, Taylor J, Tibshirani R. Supplement to “A lasso for hierarchical interactions”. 201310.1214/13-AOS1096SUPP
- Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*. 2011; 3:1–124.
- Breiman L. Better subset regression using the nonnegative garrote. *Technometrics*. 1995; 37:373–384.
- Breiman, L.; Friedman, JH.; Olshen, RA.; Stone, CJ. *Classification and Regression Trees*. Wadsworth Advanced Books and Software; Belmont, CA: 1984.
- Chipman H. Bayesian variable selection with related predictors. *Canad J Statist*. 1996; 24:17–36.
- Choi NH, Li W, Zhu J. Variable selection with the strong heredity constraint and its oracle property. *J Amer Statist Assoc*. 2010; 105:354–364.
- Cox DR. *Interaction*. *Internat Statist Rev*. 1984; 52:1–31.

- Efron B. How biased is the apparent error rate of a prediction rule? *J Amer Statist Assoc.* 1986; 81:461–470.
- Efron B, Hastie T, Johnstone I, Tibshirani R. Least angle regression. *Ann Statist.* 2004; 32:407–499.
- Forina, M.; Armanino, C.; Lanteri, S.; Tiscornia, E. *Food Research and Data Analysis*. Applied Science Publishers; London: 1983. Classification of olive oils from their fatty acid composition; p. 189-214.
- Friedman JH. Multivariate adaptive regression splines (with discussion). *Ann Statist.* 1991; 19:1–141.
- Friedman JH, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software.* 2010; 33:1–22. [PubMed: 20808728]
- George E, McCulloch R. Variable selection via gibbs sampling. *J Amer Statist Assoc.* 1993; 88:884–889.
- Hamada M, Wu C. Analysis of designed experiments with complex aliasing. *Journal of Quality Technology.* 1992; 24:130–137.
- Jeantton R, Audibert J-Y, Bach F. Structured variable selection with sparsity-inducing norms. *J Mach Learn Res.* 2011; 12:2777–2824.
- Jeantton R, Mairal J, Obozinski G, Bach F. Proximal methods for sparse hierarchical dictionary learning. In: *Proceedings of the International Conference on Machine Learning (ICML)*. 2010
- McCullagh, P.; Nelder, JA. *Generalized Linear Models*. Chapman & Hall; London: 1983.
- Nardi Y, Rinaldo A. The log-linear group-lasso estimator and its asymptotic properties. *Bernoulli.* 2012; 18:945–974.
- Nelder JA. A reformulation of linear models. *J Roy Statist Soc Ser A.* 1977; 140:48–76.
- Nelder JA. Letters to the editors: Functional marginality is important. *J R Stat Soc Ser C Appl Stat.* 1997; 46:281–286.
- Obozinski G, Jacob L, Vert J. Group lasso with overlaps: The latent group lasso approach. 2011 Available at arXiv:1110.0413.
- Park M, Hastie T. Penalized logistic regression for detecting gene interactions. *Biostatistics.* 2008; 9:30–50. [PubMed: 17429103]
- Peixoto J. Hierarchical variable selection in polynomial regression models. *Amer Statist.* 1987; 41:311–313.
- Radchenko P, James GM. Variable selection using adaptive nonlinear interaction structures in high dimensions. *J Amer Statist Assoc.* 2010; 105:1541–1553.
- Rhee S, Taylor J, Wadhera G, Ben-Hur A, Brutlag D, Shafer R. Genotypic predictors of human immunodeficiency virus type 1 drug resistance. *Proc Natl Acad Sci USA.* 2006; 103:17355. [PubMed: 17065321]
- Stein CM. Estimation of the mean of a multivariate normal distribution. *Ann Statist.* 1981; 9:1135–1151.
- Tibshirani R. Regression shrinkage and selection via the lasso. *J Roy Statist Soc Ser B.* 1996; 58:267–288.
- Tibshirani RJ, Taylor J. The solution path of the generalized lasso. *Ann Statist.* 2011; 39:1335–1371.
- Tibshirani RJ, Taylor J. Degrees of freedom in lasso problems. *Ann Statist.* 2012; 40:1198–1232.
- Tseng P. Convergence of a block coordinate descent method for nondifferentiable minimization. *J Optim Theory Appl.* 2001; 109:475–494.
- Turlach B. Discussion of “Least angle regression”. *Ann Statist.* 2004; 32:481–490.
- Wu J, Devlin B, Ringquist S, Trucco M, Roeder K. Screen and clean: A tool for identifying interactions in genome-wide association studies. *Genetic Epidemiology.* 2010; 34:275–285. [PubMed: 20088021]
- Yuan M, Joseph VR, Lin Y. An efficient variable selection approach for analyzing designed experiments. *Technometrics.* 2007; 49:430–439.
- Yuan M, Joseph VR, Zou H. Structured variable selection and estimation. *Ann Appl Stat.* 2009; 3:1738–1757.
- Yuan M, Lin Y. Model selection and estimation in regression with grouped variables. *J R Stat Soc Ser B Stat Methodol.* 2006; 68:49–67.

Zhao P, Rocha G, Yu B. The composite absolute penalties family for grouped and hierarchical variable selection. *Ann Statist.* 2009; 37:3468–3497.

Zou H, Hastie T. Regularization and variable selection via the elastic net. *J R Stat Soc Ser B Stat Methodol.* 2005; 67:301–320.

Zou H, Hastie T, Tibshirani R. On the “degrees of freedom” of the lasso. *Ann Statist.* 2007; 35:2173–2192.

APPENDIX A: PROOFS OF STRONG AND WEAK HIERARCHY

We begin by proving Lemma 1, which characterizes all solutions to (5) as a relatively simple function of y . The structure of our proof is based on (Ryan) Tibshirani and Taylor (2011, 2012).

A.1. Characterizing the solution

For ease of analysis, we write (5) equivalently in terms of Θ^+ and Θ^- . Also, for notational simplicity, we write $\varphi = (\beta^+, \beta^-, \Theta^+, \Theta^-)$ and $X = (X; -X; Z/2; -Z/2)$. The strong hierarchical lasso problem is the following:

$$\begin{aligned} & \text{Minimize}_{\beta^+, \beta^-, \Theta^+, \Theta^-} \frac{1}{2} \|y - \tilde{X}\phi\|^2 + \lambda_1 1^T(\beta^+ + \beta^-) + \lambda_2 \langle 11^T, \Theta^+ + \Theta^- \rangle \\ \text{s.t.} \quad & 1^T(\Theta_j^+ + \Theta_j^-) \leq \beta_j^+ + \beta_j^- \text{ and } \beta_j^+ \geq 0, \beta_j^- \geq 0 \text{ for each } j, \Theta^+ - \Theta^- = \Theta^{+T} - \Theta^{-T}, \Theta_{jk}^\pm \geq 0, \Theta_{jj}^\pm = 0. \end{aligned}$$

In introducing Θ^\pm , we are not in fact changing the problem since at a solution $\hat{\Theta}^\pm = \max\{\pm\Theta, 0\}$ (for $\lambda_2 > 0$). To see this, note that given any feasible point with $\Theta_{jk}^+ > 0$ and $\Theta_{jk}^- > 0$, we can produce a feasible point with strictly lower objective by reducing $\Theta_{jk}^+, \Theta_{jk}^-, \Theta_{kj}^+, \Theta_{kj}^-$ all by equal amounts.

We will try to make this as close as possible in form and notation to (Ryan) Tibshirani and Taylor’s (2012) treatment of the generalized lasso problem. Our optimization problem is of the form

$$\text{Minimize}_{\phi} \frac{1}{2} \|y - \tilde{X}\phi\|^2 + w^T \phi \quad \text{s.t.} \quad D\phi \geq 0, L\phi = 0. \quad (13)$$

The Lagrangian of this problem is

$$L(\phi; \mu, \nu) = \frac{1}{2} \|y - \tilde{X}\phi\|^2 + w^T \phi - \mu^T D\phi + \nu^T L\phi,$$

where $\mu \geq 0$ and ν are dual variables. The KKT conditions for $(\hat{\phi}(y), (\hat{\mu}(y), \hat{\nu}(y)))$ to be an optimal primal-dual pair are the following:

$$\begin{aligned} \tilde{X}^T (y - \tilde{X}\hat{\phi}) &= w - D^T \hat{\mu} + L^T \hat{\nu}, \\ \hat{\mu}_i (D\hat{\phi})_i &= 0, \\ \hat{\mu} &\geq 0, \\ D\hat{\phi} &\geq 0, \quad L\hat{\phi} = 0. \end{aligned}$$

Now, define the “boundary” and “active” sets as

$$\begin{aligned} \mathcal{B}(\hat{\mu}) &= \{i: \hat{\mu}_i = 0\}, \\ \mathcal{A}(\hat{\phi}) &= \{i: [D\hat{\phi}]_i > 0\}. \end{aligned}$$

These are not necessarily unique since $(\varphi, (\hat{\mu}, \hat{\nu}))$ may not be unique. In terms of the active set $\mathcal{A}(\hat{\phi})$, the KKT conditions become

$$\begin{aligned} \tilde{X}^T (y - \tilde{X}\hat{\phi}) &= w - D_{-\mathcal{A}(\hat{\phi})}^T \hat{\mu}_{-\mathcal{A}(\hat{\phi})} + L^T \hat{\nu}, \\ \hat{\mu}_{\mathcal{A}(\hat{\phi})} &= 0, \quad \hat{\mu} \geq 0, \\ L\hat{\phi} &= 0, \quad D_{-\mathcal{A}(\hat{\phi})}\hat{\phi} = 0. \end{aligned}$$

Solving for $\hat{\phi}$, we get the following characterization of a strong hierarchical lasso solution:

Lemma 1

Suppose $\hat{\phi}$ is a solution to the strong hierarchical lasso problem (5) [taking $q(\beta_0, \beta, \Theta)$ to be the quadratic loss] with $\mathcal{A}(\hat{\phi}) = \{i: [D\hat{\phi}]_i > 0\}$. Then, $\hat{\phi}$ can be written in terms of $\mathcal{A}(\hat{\phi})$ and y as

$$\hat{\phi} = (\tilde{X}P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})})^+ (y - (P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})})\tilde{X}^T)^+ w) + b,$$

where $b \in \text{null}(X) \cap \text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})$ satisfies

$$D_i[(\tilde{X}P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})})^+ (y - (P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})})\tilde{X}^T)^+ w) + b] > 0$$

for all $i \in \mathcal{A}(\hat{\phi})$.

Proof—Defining $\tilde{D} = \begin{pmatrix} D_{-\mathcal{A}(\hat{\phi})} \\ L \end{pmatrix}$ and $P = P_{\text{null}(\tilde{D})} = P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})}$, we solve for $\hat{\phi}$ in the same manner as is done in (Ryan) Tibshirani and Taylor (2012). Since $D\hat{\phi} = 0$ is equivalent to $P\hat{\phi} = 0$, we have $PX^T(y - X\hat{\phi}) = Pw$. We see that $Pw \in \text{col}(PX^T)$ and thus $Pw = (PX^T)^+(PX^T)^+Pw$. Thus, $PX^T\tilde{X}\hat{\phi} = PX^T(y - (PX^T)^+Pw)$ from which we get

$$\hat{\phi} = (\tilde{X}P)^+ (y - (P\tilde{X}^T)^+ Pw) + b$$

for $b \in \text{null}(XP)$ and such that $D\tilde{b} = 0$ and $D_i\hat{\phi} > 0$ for $i \in \mathcal{A}(\hat{\phi})$. To complete the result, we observe that the first two conditions reduce to $b \in \text{null}(X) \cap \text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})$.

We will use this characterization of a solution both to prove that the hierarchy property holds with probability one under weak assumptions and to derive an unbiased estimate of the degrees of freedom.

Before we do so, we write out $\tilde{D} = \begin{pmatrix} D_{-\mathcal{A}(\hat{\phi})} \\ L \end{pmatrix}$ more explicitly and introduce a little notation that will be useful later. Every row of D corresponds to an inequality constraint, and we can describe these rows in terms of ten subsets,

$$\begin{aligned} \mathcal{L} &= \{j: 1^T(\hat{\Theta}_j^+ + \hat{\Theta}_j^-) < \hat{\beta}_j^+ + \hat{\beta}_j^-\}, & \mathcal{T} &= \mathcal{L}^c, \\ \mathcal{P}(\hat{\beta}^\pm) &= \{j: \hat{\beta}_j^\pm > 0\}, & \mathcal{Z}(\hat{\beta}^\pm) &= \mathcal{P}(\hat{\beta}^\pm)^c, \\ \mathcal{P}(\hat{\Theta}^\pm) &= \{j \neq k: \hat{\Theta}_{jk}^\pm > 0\}, & \mathcal{Z}(\hat{\Theta}^\pm) &= \mathcal{P}(\hat{\Theta}^\pm)^c. \end{aligned} \tag{14}$$

The set $\mathcal{A}(\hat{\phi})^c$ is made up of \mathcal{T} , $\mathcal{Z}(\hat{\beta}^+)$, $\mathcal{Z}(\hat{\beta}^-)$, $\mathcal{Z}(\hat{\Theta}^+)$ and $\mathcal{Z}(\hat{\Theta}^-)$. The matrix \tilde{D} has $2p+2p^6$ columns that can be partitioned as $(D^{\beta^+} : D^{\beta^-} : D^{\Theta^+} : D^{\Theta^-})$ and a row for every constraint. The rows of this matrix are the following (where e_j and 1_p are row vectors):

- R1. For each $j \in \mathcal{T}$, $(e_j \quad e_j \quad -e_j \otimes 1_p \quad -e_j \otimes 1_p \quad)$
- R2. For each $j \in \mathcal{Z}(\beta^+)$, $(e_j \quad 0 \quad 0 \quad 0 \quad)$
- R3. For each $j \in \mathcal{Z}(\beta^-)$, $(0 \quad e_j \quad 0 \quad 0 \quad)$
- R4. For each $jk \in \mathcal{Z}(\Theta^+)$, $(0 \quad 0 \quad e_j \otimes e_k \quad 0 \quad)$
- R5. For each $jk \in \mathcal{Z}(\Theta^-)$, $(0 \quad 0 \quad 0 \quad e_j \otimes e_k \quad)$
- R6. For each j , $(0 \quad 0 \quad e_j \otimes e_j \quad 0 \quad)$
- R7. For each j , $(0 \quad 0 \quad 0 \quad e_j \otimes e_j \quad)$
- R8. For each $j < k$, $(0 \quad 0 \quad e_j \otimes e_k + e_k \otimes e_j \quad -e_j \otimes e_k - e_k \otimes e_j \quad)$.

We will refer to this in the proofs that follow.

A.2. Proof of strong hierarchy (Property 2)

Including the elastic net penalty, $(\varepsilon/2)\|\beta^+\|^2 + (\varepsilon/2)\|\beta^-\|^2 + (\varepsilon/2)\|\Theta\|_F^2$, is equivalent to replacing X and y in (13) by

$$\tilde{X}_\varepsilon = \begin{pmatrix} X & -X & Z/2 & -Z/2 \\ \sqrt{\varepsilon}I_p & 0 & 0 & 0 \\ 0 & \sqrt{\varepsilon}I_p & 0 & 0 \\ 0 & 0 & \sqrt{\varepsilon}I_{p^2} & -\sqrt{\varepsilon}I_{p^2} \end{pmatrix} \quad \text{and} \quad y_\varepsilon = \begin{pmatrix} y \\ 0_{(2p+p^2) \times 1} \end{pmatrix}.$$

Suppose we solve (13) with the above design matrix. By Lemma 1,

$$\hat{\phi} = (\tilde{X}_\varepsilon P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})})^+ (y_\varepsilon - (P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})}) \tilde{X}_\varepsilon^T)^+ w) + b$$

for some $b \in \text{null}(X_\varepsilon) \cap \text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})$ satisfying

$$D_i[(\tilde{X}_\varepsilon P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})})^+ (y_\varepsilon - (P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})}) \tilde{X}_\varepsilon^T)^+ w) + b] > 0$$

for all $i \in \mathcal{A}(\hat{\phi})$. Let $S_v : \mathbb{R}^{2p+2p^2} \rightarrow \mathbb{R}^{|v|}$ be the linear operator that selects the part of a vector corresponding to the variable v . Now, $b \in \text{null}(X_\varepsilon)$ implies that $S_{\beta^+}(b) = S_{\beta^-}(b) = 0$ and $S_{\Theta^+}(b) = -S_{\Theta^-}(b)$. We showed earlier that we cannot have $\hat{\Theta}_{jk}^+ > 0$ and $\hat{\Theta}_{jk}^- > 0$. This means that for any jk , there must be an $i \notin \mathcal{A}(\hat{\phi})$ for which $D_i \hat{\phi} = 0$ corresponds to $S_{\Theta_{jk}^+}(\hat{\phi}) = 0$ or $S_{\Theta_{jk}^-}(\hat{\phi}) = 0$. Thus, $D_{-\mathcal{A}(\hat{\phi})} b = 0$ means that $S_{\Theta_{jk}^+}(b) = 0$ or $S_{\Theta_{jk}^-}(b) = 0$ for each jk . This implies that $\text{null}(X_\varepsilon) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})}) = \{0\}$ and thus $b = 0$.

We show now that $P(\cup_j \{\hat{\beta}_j^+ = \hat{\beta}_j^- > 0\}) = 0$. In terms of our above notation, this is

$$P\left(\cup_j \{\hat{\phi}^T(e_j^+ - e_j^-) = 0, \hat{\phi}^T e_j^+ > 0, \hat{\phi}^T e_j^- > 0\}\right) = 0, \quad (15)$$

where $e_j^\pm \in \mathbb{R}^{2p+2p^2}$ is the vector with all zeros except for $S_{\beta^\pm}(e_j^\pm) = 1$. Let $P = P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}})}$, and consider the set

$$\mathcal{N} = \cup_{\mathcal{A}} \cup_{j=1}^p \{z : [(\tilde{X}_\varepsilon P)^+ (z - (P \tilde{X}_\varepsilon^T)^+ w)]^T (e_j^+ - e_j^-) = 0, [(\tilde{X}_\varepsilon P)^+ (z - (P \tilde{X}_\varepsilon^T)^+ w)]^T e_j^+ > 0, [(\tilde{X}_\varepsilon P)^+ (z - (P \tilde{X}_\varepsilon^T)^+ w)]^T e_j^- > 0\}.$$

In light of (14), fixing \mathcal{A} automatically specifies \mathcal{T} , $\mathcal{P}(\hat{\beta}^\pm)$, $\mathcal{P}(\hat{\Theta}^\pm)$. The outer union is restricted to those subsets \mathcal{A} of $2p + 2p^2$ elements that would have $\mathcal{P}(\hat{\Theta}^+) \cap \mathcal{P}(\hat{\Theta}^-) = \emptyset$ and $\mathcal{P}(\hat{\beta}^+) \cap \mathcal{P}(\hat{\beta}^-) \subseteq \mathcal{T}$. The event in (15) is contained in $\{y_\varepsilon \in \mathcal{N}\}$ since it corresponds to the case in which \mathcal{A} is $\mathcal{A}(\hat{\phi})$. We begin by showing that $\mathcal{A}(\hat{\phi})$ is in this restricted union with probability one. We have already argued that at a solution we must have $\hat{\Theta}_{jk}^+ \hat{\Theta}_{jk}^- = 0$ for all jk . Now, $\hat{\beta}_j^+ > 0$ and $\hat{\beta}_j^- > 0$ together imply that $1^T(\hat{\Theta}_j^+ + \hat{\Theta}_j^-) = \hat{\beta}_j^+ + \hat{\beta}_j^-$ since otherwise we could lower the objective by reducing $\hat{\beta}_j^+$ and $\hat{\beta}_j^-$ without leaving the feasible set. Therefore, it would be sufficient to show that $P(y_\varepsilon \in \mathcal{N}) = 0$. We do so by observing that $\{y_\varepsilon \in \mathcal{N}\}$ is a finite union of zero probability sets.

We begin by establishing that $\text{null}(I_{1:n}(P\tilde{X}_\varepsilon^T)^+) = \text{null}(\tilde{X}P)$. To do so, we write $P = UU^T$ for some $U^T U = I$. Now, $\text{row}(\tilde{X}P) \subseteq \text{row}(P) = \text{col}(U)$, so we can write $U = (U_1 : U_2)$, where $\text{row}(\tilde{X}P) = \text{col}(U_1)$ and thus $\tilde{X}P U_2 = 0$. Since $P U_2 = U_2$, it follows that $\tilde{X} U_2 = 0$. Write

$$U_i^T = (U_i^{\beta+T} : U_i^{\beta-T} : U_i^{\Theta+T} : U_i^{\Theta-T}) \text{ for } i = 1, 2, \text{ and observe that}$$

$$U_i^T \tilde{X}_\varepsilon^T = [U_i^T \tilde{X}^T : \sqrt{\varepsilon} U_i^{\beta+T} : \sqrt{\varepsilon} U_i^{\beta-T} : \sqrt{\varepsilon} (U_i^{\Theta+T} - U_i^{\Theta-T})] \text{ so that}$$

$$\begin{aligned} U_1^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_2 &= 0 + \varepsilon [U_1^{\beta+T} U_2^{\beta+} + U_1^{\beta-T} U_2^{\beta-} + (U_1^{\Theta+} - U_1^{\Theta-}) (U_2^{\Theta+} - U_2^{\Theta-})] \\ &= \varepsilon [U_1^T U_2 - U_1^{\Theta+T} U_2^{\Theta-} - U_1^{\Theta-T} U_2^{\Theta+}] \\ &= -\varepsilon \left[0 + \sum_{jk} [U_1^{\Theta+}]_{jk} [U_2^{\Theta-}]_{jk} + [U_1^{\Theta-}]_{jk} [U_2^{\Theta+}]_{jk} \right]. \end{aligned}$$

Now, for each jk we must have $jk \in \mathcal{Z}(\Theta^+) \cup \mathcal{Z}(\Theta^-)$ since $\mathcal{P}(\Theta^+) \cap \mathcal{P}(\Theta^-) = \emptyset$. Thus, for each jk , there is a R4 or R5 row in D_{-A} and thus $D_{-A} (U_1 : U_2) = 0$ implies that $[U_1^{\Theta+}]_{jk} = 0$ or $[U_2^{\Theta-}]_{jk} = 0$ for each jk and likewise $[U_2^{\Theta+}]_{jk} = 0$ or $[U_1^{\Theta-}]_{jk} = 0$. Therefore, $U_1^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_2 = 0$. Now,

$$\begin{aligned} I_{1:n}(P\tilde{X}_\varepsilon^T)^+ &= \tilde{X}P(P\tilde{X}_\varepsilon^T \tilde{X}_\varepsilon P)^+ = \tilde{X}U U^T (U U^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U U^T)^+ \\ &= \tilde{X}U U^T U (U^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U)^+ U^T = \tilde{X}U (U^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U)^+ U^T \\ &= \tilde{X} (U_1 U_2) \begin{pmatrix} U_1^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_1 & U_1^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_2 \\ U_2^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_1 & U_2^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_2 \end{pmatrix}^+ U^T \\ &= \tilde{X} (U_1 0) \begin{pmatrix} U_1^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_1 & 0 \\ 0 & U_2^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_2 \end{pmatrix}^+ U^T \\ &= \tilde{X} U_1 (U_1^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_1)^+ U_1^T. \end{aligned}$$

Now, $U_1^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_1 \succ 0$ since $U_1^T \tilde{X}_\varepsilon^T \tilde{X}_\varepsilon U_1 = U_1^T \tilde{X}^T \tilde{X} U_1 + J$ and $\succeq 0$ and $\tilde{X} U_1$ has full column rank. Thus, $\text{null}(I_{1:n}(P\tilde{X}_\varepsilon^T)^+) = \text{null}(U_1^T) = \text{null}(\tilde{X}P)$. This completes the first part of the proof.

Next, we show that $\tilde{X}P(e_j^+ - e_j^-) \neq 0$ as long as $[(\tilde{X}_\varepsilon P)^+(y_\varepsilon - (P\tilde{X}_\varepsilon^T)^+ w)] \times e_j^\pm > 0$. Now, since $j \in \mathcal{P}(\beta^+) \cap \mathcal{P}(\beta^-) \subseteq \mathcal{T}$, the only row of D that has $\tilde{D}_i(e_j^\pm) \neq 0$ is the R1 row; but clearly $\tilde{D}_i(e_j^+ - e_j^-) = 1 - 1 = 0$. Thus, $e_j^+ - e_j^- \in \text{null}(\tilde{D})$ and $P(e_j^+ - e_j^-) = e_j^+ - e_j^-$. It follows that

$$\tilde{X}P(e_j^+ - e_j^-) = \begin{pmatrix} 2x_j \\ \sqrt{\varepsilon} e_j \\ -\sqrt{\varepsilon} e_j \\ 0 \end{pmatrix} \neq 0 \quad \text{assuming } \varepsilon > 0.$$

Putting these two parts of the proof together establishes that $I_{1:n}(P\tilde{X}_\varepsilon^T)^+(e_j^+ - e_j^-) \neq 0$. Thus, $\{y_\varepsilon \in \mathcal{N}\}$ is a finite union of Lebesgue measure 0 sets. This shows that $P(y_\varepsilon \in \mathcal{N}) = 0$ as long as y is absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^n .

A.3. Proof of weak hierarchy (Property 3)

An argument nearly identical to that of the previous section establishes that

$\hat{\Theta}_{jk} \neq 0 \Rightarrow \hat{\beta}_j^+ - \hat{\beta}_j^- \neq 0$ with probability one. Thus, if $\hat{\beta}_j^+ - \hat{\beta}_j^- = 0$ and $\hat{\beta}_k^+ - \hat{\beta}_k^- = 0$, then both $\hat{\Theta}_{jk} = 0$ and $\hat{\Theta}_{kj} = 0$. It follows then that $(\hat{\Theta}_{jk} + \hat{\Theta}_{kj})/2 = 0$. This establishes weak hierarchy.

APPENDIX B: DEGREES OF FREEDOM

B.1. Proof of unbiased estimate (Property 4)

The fit in terms of the active set is given by

$$\tilde{X}\hat{\phi} = (\tilde{X}P)(\tilde{X}P)^+(y - (P\tilde{X}^T)^+w),$$

where $P = P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})}$. Of course, $\hat{\mu}_{\mathcal{A}(\hat{\phi})} = 0$, and we can solve the KKT conditions to get the rest of the optimal dual variables in terms of the active set

$$\begin{pmatrix} \hat{\mu}_{-\mathcal{A}(\hat{\phi})} \\ \hat{\nu} \end{pmatrix} = \tilde{D}^{T+}[w - \tilde{X}^T(y - \tilde{X}\hat{\phi})] + c,$$

where $c \in \text{null}(D)$ satisfies $D^{T+}[w - X^T(y - X\hat{\phi})] + c = 0$.

Note that $\hat{\varphi} = \varphi(\hat{y})$ and thus $\mathcal{A}(\hat{\varphi})$ and b depend on y even though we do not write this explicitly. We will continue writing $\hat{\varphi}$ to mean specifically $\varphi(\hat{y})$. For y' in a neighborhood of y , we might guess that $\varphi(\hat{y}') = f(y')$ and $(\mu(\hat{y}'), \nu(\hat{y}')) = (g(y'), h(y'))$, where

$$\begin{aligned} f(y') &= (\tilde{X}P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})})^+(y' - (P_{\text{null}(L) \cap \text{null}(D_{-\mathcal{A}(\hat{\phi})})})\tilde{X}^T)^+w) + b, \\ \begin{pmatrix} g(y') \\ h(y') \end{pmatrix} &= \begin{pmatrix} D_{-\mathcal{A}(\hat{\phi})} \\ L \end{pmatrix}^{T+} [w - \tilde{X}^T(y' - \tilde{X}f(y'))] + c, \\ g(y')_{\mathcal{A}(\hat{\phi})} &= 0. \end{aligned}$$

To verify this guess, we need to check that the pair $(f(y'), (g(y'), h(y')))$ satisfies the optimality conditions at y' .

$$\begin{aligned} \tilde{X}^T (y' - \tilde{X}f(y')) &= w - D^T g(y') + L^T h(y'), \\ g_i(y') (Df(y'))_i &= 0, \\ g(y') \geq 0, \quad Df(y') \geq 0, \quad Lf(y') &= 0. \end{aligned}$$

First of all, $Lf(y') = 0$ holds since $L(X_{P_{\text{null}(L)/\text{null}(D_{-\mathcal{A}(\hat{\phi})})}^+})^+ = 0$ and $Lb = 0$. Likewise, $D_{-\mathcal{A}(\hat{\phi})} f(y') = 0$. Now $D_{\mathcal{A}(\hat{\phi})} f(y') > 0$, so by continuity of f , we have $D_{\mathcal{A}(\hat{\phi})} f(y') > 0$ for all y' in a small enough neighborhood, U_1 , of y . This establishes that $\mathcal{A}(f(y')) = \mathcal{A}(\hat{\phi})$. Now $(y')_{\mathcal{A}(\hat{\phi})} = 0$, so complementary slackness holds. To see that the first optimality condition holds, we can simply plug $(g(y'), h(y'))$ into the left-hand side. All that remains is to show that $(y')_{-\mathcal{A}(\hat{\phi})} = 0$. If we knew that $\hat{\mu}_{-\mathcal{A}(\hat{\phi})} > 0$, then by continuity of g we could argue that over a small enough neighborhood, U_2 , $(y')_{-\mathcal{A}(\hat{\phi})} > 0$. However, it could be the case that $\mu_i = 0$ for some $i \notin \mathcal{A}(\hat{\phi})$, that is, $i \in \mathcal{B}(\hat{\phi}) \setminus \mathcal{A}(\hat{\phi})$. Nonetheless, one can show that there is a set \mathcal{N} of measure 0 for which $y \notin \mathcal{N}$ implies that $\mathcal{A}(\hat{\phi}(y)) = \mathcal{A}(\hat{\phi}(y'))$ and $\mathcal{B}(\hat{\phi}(y)) = \mathcal{B}(\hat{\phi}(y'))$ for all y' in a neighborhood of y . Lemma 9 of (Ryan Tibshirani and Taylor (2012)) proves this result for a nearly identical situation.

The fit $X\hat{\phi}(y)$ is a piecewise affine function of y . Using Stein's formula for the degrees of freedom [as described in Ryan, Tibshirani and Taylor (2012)], we get that

$$df(\tilde{X}\hat{\phi}) = E[\nabla \cdot \tilde{X}\hat{\phi}(y)] = E[\text{tr}\{(\tilde{X}P)(\tilde{X}P)^+\}] = E[\text{rank}(\tilde{X}P)],$$

where $P = P_{\text{null}(L)/\text{null}(D_{-\mathcal{A}(\hat{\phi})})}$.

B.2. Proof of bound on estimate (Property 5)

We bound this by an estimate that is more interpretable:

$$\text{rank}(\tilde{X}P) \leq \text{rank}(P) = \text{nullity} \begin{pmatrix} L \\ D_{-\mathcal{A}(\hat{\phi})} \end{pmatrix}.$$

Clearly, R2–R7 are linearly independent rows. Thus, the rank of D is at least $|\mathcal{Z}(\beta^+)| + |\mathcal{Z}(\beta^-)| + |\mathcal{Z}(\Theta^+)| + |\mathcal{Z}(\Theta^-)| + 2p$. Now, an R1 row is linearly independent of R2–R8 precisely when $j \in \mathcal{T}$ has $j \in \mathcal{Z}(\beta^+) \setminus \mathcal{Z}(\beta^-)$. To see this, note that if $j \in \mathcal{Z}(\beta^+) \setminus \mathcal{Z}(\beta^-)$, then R1 is certainly linearly independent of R2–R8 and likewise for $j \in \mathcal{Z}(\beta^-) \setminus \mathcal{Z}(\beta^+)$; however if $j \in \mathcal{Z}(\beta^+) \cap \mathcal{Z}(\beta^-)$, then $jk \in \mathcal{Z}(\Theta^+) \cap \mathcal{Z}(\Theta^-)$ for all $k \in \{1, \dots, p\} \setminus \{j\}$, and therefore this row of R1 lies in the span of R3–R7. Thus, this means there are $|\mathcal{T} \setminus (\mathcal{Z}(\beta^+) \cup \mathcal{Z}(\beta^-))|$ additional linearly independent rows. Finally, we consider R8. Clearly, R8 lies in the span of R4–R5 for $jk \in \mathcal{Z}(\Theta^+) \cap \mathcal{Z}(\Theta^-)$ since $jk \in \mathcal{Z}(\Theta^+) \implies kj \in \mathcal{Z}(\Theta^+)$ at a solution. But if $jk \in \mathcal{P}(\Theta^+) \cup \mathcal{P}(\Theta^-)$, then it is linearly independent of R1–R8. Therefore, R8 adds $|\mathcal{P}(\Theta^+)|/2 + |\mathcal{P}(\Theta^-)|/2$ to the rank where we have used that $\mathcal{P}(\Theta^+) \cap \mathcal{P}(\Theta^-) = \emptyset$ at a solution (since $\lambda^6 > 0$) and recalling that $j < k$ for the rows of R8. In summary, we have shown that the row-rank is

$$|\mathcal{Z}(\beta^+)| + |\mathcal{Z}(\beta^-)| + |\mathcal{Z}(\Theta^+)| + |\mathcal{Z}(\Theta^-)| + 2p + |\mathcal{T} \setminus (\mathcal{Z}(\beta^+) \cup \mathcal{Z}(\beta^-))| + |\mathcal{P}(\Theta^+)|/2 + |\mathcal{P}(\Theta^-)|/2.$$

Recalling that there are $2p + 2p^6$ columns, we get that

$$\begin{aligned} \text{nullity} \begin{pmatrix} L \\ D_{-\mathcal{A}(\hat{\phi})} \end{pmatrix} &= (2p + 2p^2) - \text{rank} \begin{pmatrix} L \\ D_{-\mathcal{A}(\hat{\phi})} \end{pmatrix} \\ &= (p - |\mathcal{L}(\beta^+)|) + (p - |\mathcal{L}(\beta^-)|) + (p(p-1) - |\mathcal{L}(\Theta^+)|) + (p(p-1) - |\mathcal{L}(\Theta^-)|) - |\mathcal{T} \setminus (\mathcal{L}(\beta^+) \Delta \mathcal{L}(\beta^-))| - |\mathcal{P}(\Theta^+)|/2 - |\mathcal{P}(\Theta^-)|/2 \\ &= |\mathcal{P}(\beta^+)| + |\mathcal{P}(\beta^-)| + |\mathcal{P}(\Theta^+)|/2 + |\mathcal{P}(\Theta^-)|/2 - |\mathcal{T} \setminus (\mathcal{P}(\beta^+) \Delta \mathcal{P}(\beta^-))|. \end{aligned}$$

APPENDIX C: SOLVING THE PROX FUNCTION

The Lagrangian of (11) is given by

$$L(\beta^\pm, \theta; \alpha, \gamma^\pm) = \frac{1}{2t}(\beta^+ - \tilde{\beta}^+)^2 + \frac{1}{2t}(\beta^- - \tilde{\beta}^-)^2 + \frac{1}{2t}\|\theta - \tilde{\theta}\|^2 + \left(\frac{\lambda}{2} + \alpha\right) \|\theta\|_1 - (\gamma^+ + \alpha)\beta^+ - (\gamma^- + \alpha)\beta^-,$$

where α is the dual variable corresponding to the hierarchy constraints and γ^\pm are the dual variables corresponding to the nonnegativity constraints. For notational convenience, we have written θ for Θ_j , and we have dropped the subscripts on β^\pm . The KKT conditions are

$$\begin{aligned} (\hat{\beta}^\pm - \tilde{\beta}^\pm)/t - \hat{\gamma}^\pm - \hat{\alpha} &= 0, & (\hat{\theta} - \tilde{\theta})/t + (\lambda/2 + \hat{\alpha})u &= 0, \\ \hat{\beta}^\pm \hat{\gamma}^\pm &= 0, & \hat{\alpha}(\|\hat{\theta}\|_1 - \hat{\beta}^+ - \hat{\beta}^-) &= 0, \\ \hat{\beta}^\pm &\geq 0, & \|\hat{\theta}\|_1 &\leq \hat{\beta}^+ + \hat{\beta}^-, & \hat{\alpha} &\geq 0, & \hat{\gamma}^\pm &\geq 0, \end{aligned}$$

where u_j is a subgradient of the absolute value function evaluated at $\hat{\theta}_j$. The three conditions involving $\hat{\gamma}^\pm$ implies that $\hat{\beta}^\pm = [\tilde{\beta}^\pm + t\hat{\alpha}]_+$. The stationarity condition involving $\hat{\theta}$ implies that $\hat{\theta} = S(\tilde{\theta}, t(\lambda/2 + \hat{\alpha}))$. Now, define $f(\alpha) = \|S(\tilde{\theta}, t(\lambda/2 + \alpha))\|_1 - [\tilde{\beta}^+ + t\alpha]_+ - [\tilde{\beta}^- + t\alpha]_+$. The remaining KKT conditions involve α alone: $\alpha \hat{f}'(\alpha) = 0, \hat{f}(\alpha) \hat{\alpha} = 0, \hat{\alpha} \geq 0$. Observing that f is nonincreasing in α and piecewise linear suggests finding $\hat{\alpha}$ as done in Algorithm 3.

Algorithm 3

ONEROW: Solve (11) via dual.

Inputs: $\tilde{\beta}_j^+, \tilde{\beta}_j^- \in \mathbb{R}, \Theta_j \in \mathbb{R}^{p-1}, \lambda \geq 0$.

- 1 Find $\hat{\alpha}$. Define $f(\alpha) = \|S(\tilde{\Theta}_j, t(\lambda/2 + \alpha))\|_1 - [\tilde{\beta}_j^+ + t\alpha]_+ - [\tilde{\beta}_j^- + t\alpha]_+$.
 - a. If $f(0) = 0$, take $\hat{\alpha} = 0$ and go to step 2.
 - b. Form knot the set $\mathcal{P} = \{|\tilde{\Theta}_{jk}|/t - \lambda/2\}_{k=1}^p \cup \{-\tilde{\beta}^\pm/t\}$, and let $\mathcal{P}^+ = \mathcal{P} \cap [0, \infty)$.
 - c. Evaluate $f(p)$ for $p \in \mathcal{P}^+$.
 - d. If $f(p) = 0$ for some $p \in \mathcal{P}^+$, take $\hat{\alpha} = p$ and go to step 2.
 - e. Find adjacent knots, $p_1, p_2 \in \mathcal{P}^+$, such that $f(p_1) > 0 > f(p_2)$. Take

$$\hat{\alpha} = -f(p_1)[f(p_2) - f(p_1)] / (p_2 - p_1).$$

² Return $\hat{\Theta}_j = \mathcal{S}[\tilde{\Theta}_j, t(\lambda/2 + a)]$ and $\hat{\beta}_j^\pm = [\tilde{\beta}_j^\pm + t\hat{\alpha}]_+$.

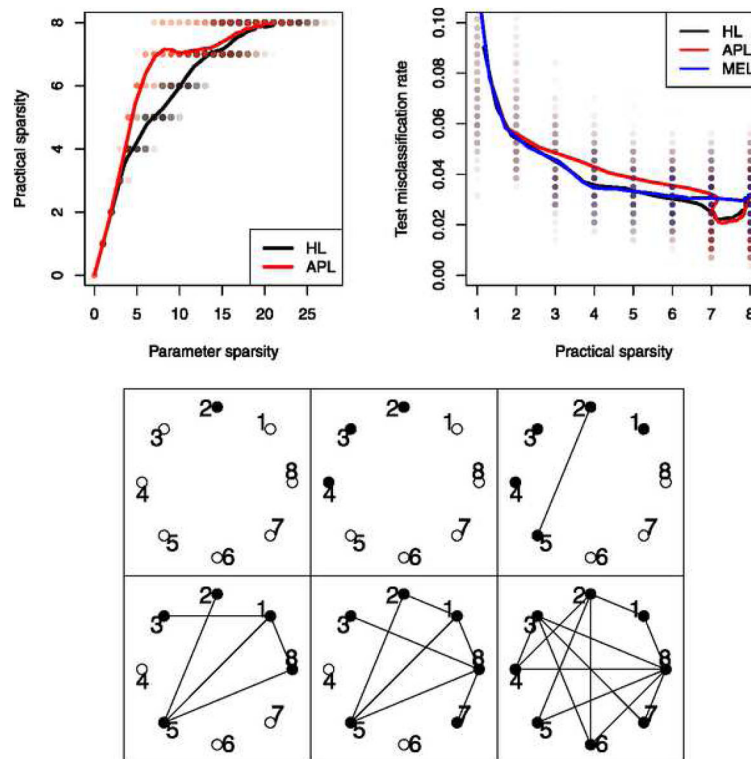


Fig. 1. Olive oil data: (Top left) Parameter sparsity is the number of nonzero coefficients while practical sparsity is the number of *measured* variables in the model. Results from all 100 random train-test splits are shown as points; lines show the average performance over all 100 runs. (Top right) Misclassification error on test set versus practical sparsity. (Bottom) Wheel plots showing the sparsity pattern at 6 values of λ for the strong hierarchical lasso. Filled nodes correspond to nonzero main effects, and edges correspond to nonzero interactions.




Fig. 2. Numerical evaluation of how well \hat{df}_λ estimates df_λ . Monte Carlo estimates of $E[\hat{df}_\lambda]$ (y-axis) versus Monte Carlo estimates of df_λ (x-axis) for a sequence of λ values (circular) are shown. One-standard-error bars are drawn and are hardly visible. Our bound on the unbiased estimate is plotted with diamonds.

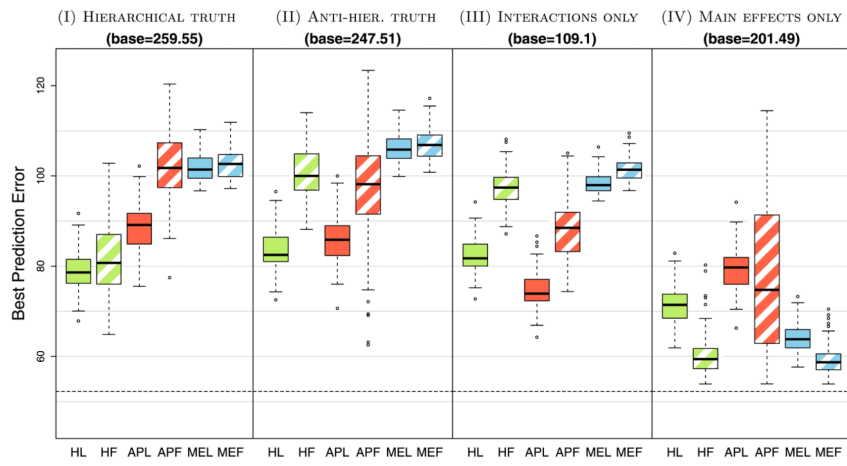


Fig. 3. Prediction error: Dashed line shows Bayes error (i.e., σ^2), and the base rate refers to the prediction error of \bar{y}_{train} . Green, red and blue colors indicate hierarchy, all-pairs, and main effect only, respectively; solid and striped indicate lasso and forward stepwise, respectively.

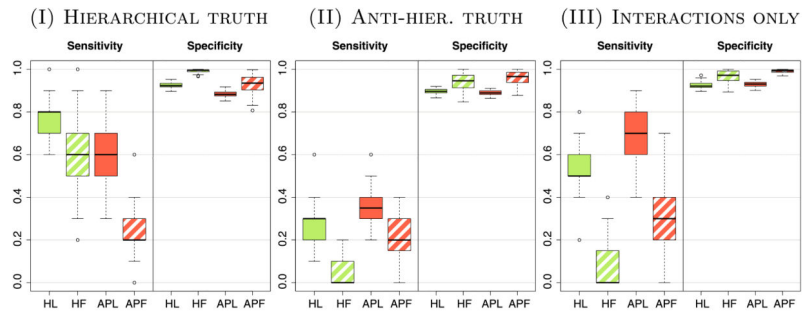


Fig. 4.

Plots show the ability of various methods to correctly recover the nonzero interactions. This is the sensitivity (i.e., proportion of $\Theta_{jk} \neq 0$ for which $\hat{\Theta}_{jk} \neq 0$) and specificity (i.e., proportion of $\Theta_{jk} = 0$ for which $\hat{\Theta}_{jk} = 0$) corresponding to the lowest prediction error model of each method.

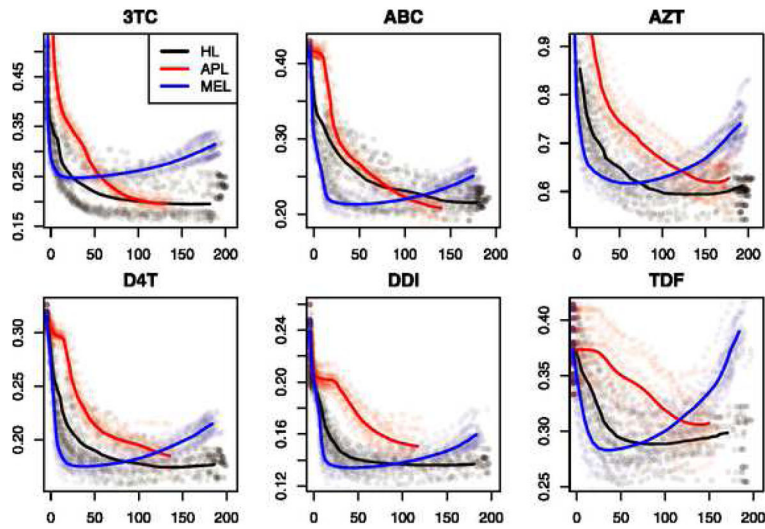


Fig. 5. HIV drug data: Test-set RMSE versus practical sparsity (i.e., number of measured variables required for prediction) for six different drugs. For each method, the data from all 20 runs are displayed in faint colors; the thick lines are averages over these runs.