

# Learning interactions via hierarchical group-lasso regularization

Michael Lim      Trevor Hastie

August 29, 2013

## Abstract

We introduce a method for learning pairwise interactions in a linear model in a manner that satisfies strong hierarchy: whenever an interaction is estimated to be nonzero, both its associated main effects are also included in the model. We motivate our approach by modeling pairwise interactions for categorical variables with arbitrary numbers of levels, and then show how we can accommodate continuous variables as well. Our approach allows us to dispense with explicitly applying constraints on the main effects and interactions for identifiability, which results in interpretable interaction models. We compare our method with existing approaches on both simulated and real data, including a genome-wide association study, all using our R package `glinternet`.

## 1 Introduction

Given an observed response and explanatory variables, we expect interactions to be present if the response cannot be explained by additive functions of the variables. The following definition makes this more precise.

**Definition 1.** *When a function  $f(x, y)$  cannot be expressed as  $g(x) + h(y)$  for some functions  $g$  and  $h$ , we say that there is an interaction in  $f$  between  $x$  and  $y$ .*

Interactions of single nucleotide polymorphisms (SNPs) are thought to play a role in cancer [Schwender and Ickstadt, 2008] and other diseases. Modeling interactions has also served the recommender systems community well: latent factor models (matrix factorization) aim to capture user-item interactions that measure a user’s affinity for a particular item, and are the

state of the art in predictive power [Koren, 2009]. In look-alike-selection, a problem that is of interest in computational advertising, one looks for features that most separates a group of observations from its complement, and it is conceivable that interactions among the features can play an important role.

Linear models scale gracefully as datasets have grown *wider* — i.e. as the number of variables increase. Problems with 10,000 or more variables are routinely handled by standard software. Scalability is a challenge when we model interactions. Even with 10,000 variables, we are already looking at a  $50 \times 10^6$ -dimensional space of possible interaction pairs. Complicating the matter are spurious correlations amongst the variables, which makes learning even harder. Finally, in some applications, sample sizes are relatively small and the signal to noise ratio is low. For example, genome wide association studies (GWAS) can involve hundreds of thousands or millions of variables, but only several thousand observations. Since the number of interactions is on the order of the square of the number of variables, computational considerations quickly become an issue.

Finding interactions is an example of the “ $p > n$ ” problem where there are more features or variables than observations. A popular approach in supervised learning problems of this type is to use regularization, such as imposing an  $\ell_2$  bound of the form  $\|\beta\|_2^2 \leq s$  (ridge) or an  $\ell_1$  bound  $\|\beta\|_1 \leq s$  (lasso) on the coefficients. The latter type of penalty has been the focus of much research since its introduction in [Tibshirani, 1996], and is called the lasso. One of the reasons for the lasso’s popularity is that it does variable selection: it sets some coefficients exactly to zero. There is a group analogue to the lasso, called the group-lasso [Yuan and Lin, 2006], that sets groups of variables to zero. The idea behind our method is to set up main effects and interactions (to be defined later) as groups of variables, and then we perform selection via the group-lasso.

Discovering interactions is an area of active research; see, for example, [Bien et al, 2013] and [Chen et al., 2011]. In this paper, we introduce **glinternet**, a method for learning first-order interactions that can be applied to categorical variables with arbitrary numbers of levels, continuous variables, and combinations of the two. Our approach consists of two phases: a screening stage (for very large problems) that gives a candidate set of main effects and interactions, followed by variable selection on the candidate set using the group-lasso. We introduce two screening procedures. The first is inspired by our observation that boosting with depth-2 trees naturally gives rise to an interaction selection process that enforces hierarchy: an interaction cannot be chosen until a split has been made on its associated

main effect. The second method is an adaptive procedure that is based on the strong rules [Tibshirani et al, 2012] for discarding predictors in lasso-type problems. We show in Section 3 how the group-lasso penalty naturally enforces strong hierarchy in the resulting solutions. While screening potentially misses some candidate interactions, we note that our approach can handle very large problems without any screening. In Section 7.4 we fit our model with 27K Snps (three-level factors), 3,500 observations, and 360M candidate interactions, without any screening.

To summarize, our method has consists of two phases:

1. If required, screen the variables to get a candidate set  $\mathcal{C}$  of interactions and their associated main effects. Otherwise, take  $\mathcal{C}$  to consist of all main effects and pairwise interactions.
2. Fit a suitable group-lasso model using the candidate terms in  $\mathcal{C}$ , and with a grid of values for the regularization parameter. Start with  $\lambda = \lambda_{max}$  for which all terms are zero. As  $\lambda$  is decreased, more terms enter the model, and we stop once a user-specified number of interactions have been discovered. Alternatively, we can choose  $\lambda$  using any model selection technique such as cross validation.

## 1.1 A simulated example

As a first example and a forward peek at the results, we perform 100 simulations with 500 3-level categorical variables and 800 observations (squared error loss with quantitative response). There are 10 main effects and 10 interactions in the ground truth, and the noise level is chosen to give a signal to noise ratio of one. We run `glinternet` without any screening, and stop after ten interactions have been found. The average false discovery rate and standard errors are plotted as a function of the number of interactions found in Figure 1.

## 1.2 Organization of the paper

The rest of the paper is organized as follows. Section 2 introduces the problem and notation. In Section 3, we introduce the group-lasso and how it fits into our framework for finding interactions. We also show how `glinternet` is equivalent to an overlapped group lasso. We discuss screening in Section 4, and give several examples with both synthetic and real datasets in




Figure 1: False discovery rate vs number of discovered interactions, averaged over 100 simulations with 500 3-level factors. The ground truth has 10 main effects and 10 interactions.

Section 7 before going into algorithmic details in Section 8. We conclude with a discussion in Section 9.

## 2 Background and notation

We use the random variables  $Y$  to denote the response,  $F$  to denote a categorical feature, and  $Z$  to denote a continuous feature. We use  $L$  to denote the number of levels that  $F$  can take. For simplicity of notation we will use the first  $L$  positive integers to represent these  $L$  levels, so that  $F$  takes values in the set  $\{i \in \mathbb{Z} : 1 \leq i \leq L\}$ . Each categorical variable has an associated random dummy vector variable  $X \in \mathbb{R}^L$  with a 1 that indicates which level  $F$  takes, and 0 everywhere else.

When there are  $p$  categorical (or continuous) features, we will use subscripts to index them, i.e.  $F_1, \dots, F_p$ . Boldface font will always be reserved for vectors or matrices that comprise of realizations of these random variables. For example,  $\mathbf{Y}$  is the  $n$ -vector of observations of the random variable  $Y$ , likewise for  $\mathbf{F}$  and  $\mathbf{Z}$ . Similarly,  $\mathbf{X}$  is a  $n \times L$  indicator matrix whose  $i$ -th

row consists of a 1 in the  $F_i$ -th column and 0 everywhere else. We use a  $n \times (L_i \cdot L_j)$  indicator matrix  $\mathbf{X}_{i:j}$  to represent the interaction  $F_i : F_j$ . We will write

$$\mathbf{X}_{i:j} = \mathbf{X}_i * \mathbf{X}_j \quad (1)$$

for the matrix consisting of all pairwise products of columns of  $\mathbf{X}_i$  and  $\mathbf{X}_j$ . For example,

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} * \begin{pmatrix} e & f \\ g & h \end{pmatrix} = \begin{pmatrix} ae & af & be & bf \\ cg & ch & dg & dh \end{pmatrix}. \quad (2)$$

## 2.1 Definition of interaction for categorical variables

To see how Definition 1 applies to this setting, let  $\mathbb{E}(Y|F_1 = i, F_2 = j) = \mu_{ij}$ , the conditional mean of  $Y$  given that  $F_1$  takes level  $i$ , and  $F_2$  takes level  $j$ . There are 4 possible cases:

1.  $\mu_{ij} = \mu$  (no main effects, no interactions)
2.  $\mu_{ij} = \mu + \theta_1^i$  (one main effect  $F_1$ )
3.  $\mu_{ij} = \mu + \theta_1^i + \theta_2^j$  (two main effects)
4.  $\mu_{ij} = \mu + \theta_1^i + \theta_2^j + \theta_{1:2}^{ij}$  (main effects and interaction)

Note that all but the first case is overparametrized, and a common remedy is to impose sum constraints on the main effects and interactions:

$$\sum_{i=1}^{L_1} \theta_1^i = 0, \quad \sum_{j=1}^{L_2} \theta_2^j = 0 \quad (3)$$

and

$$\sum_{i=1}^{L_1} \theta_{1:2}^{ij} = 0 \text{ for fixed } j, \quad \sum_{j=1}^{L_2} \theta_{1:2}^{ij} = 0 \text{ for fixed } i. \quad (4)$$

In what follows,  $\theta_i$ ,  $i = 1, \dots, p$ , will represent the main effect coefficients, and  $\theta_{i:j}$  will denote the interaction coefficients. We will use the terms “main effect coefficients” and “main effects” interchangeably, and likewise for interactions.

## 2.2 Weak and strong hierarchy

An interaction model is said to obey strong hierarchy if an interaction can be present only if both of its main effects are present. Weak hierarchy is obeyed as long as either of its main effects are present. Since main effects as defined above can be viewed as deviations from the global mean, and interactions are deviations from the main effects, it rarely make sense to have interactions without main effects. This leads us to prefer interaction models that are hierarchical. We will see in Section 3 that `glinternet` produces estimates that obey strong hierarchy.

## 2.3 First order interaction model

Our model for a quantitative response  $Y$  is given by

$$E(Y|X) = \mu + \sum_{i=1}^p X_i \theta_i + \sum_{i < j} X_{i:j} \theta_{i:j}, \quad (5)$$

and for a binary response, we have

$$\text{logit}(\mathbb{P}(Y = 1|X)) = \mu + \sum_{i=1}^p X_i \theta_i + \sum_{i < j} X_{i:j} \theta_{i:j}. \quad (6)$$

We fit these models by minimizing an appropriate choice of loss function  $\mathcal{L}$ . Because the models are still overparametrized, we impose the relevant constraints for the coefficients  $\theta$  (see (3) and (4)). We can thus cast the problem of fitting a first-order interaction model as an optimization problem with constraints:

$$\text{argmin}_{\mu, \theta} \mathcal{L}(\mathbf{Y}; \mu, \theta) \quad (7)$$

subject to the relevant constraints.  $\mathcal{L}$  can be any loss function, typically squared error loss for the quantitative response:

$$\mathcal{L}(\mathbf{Y}; \mu, \theta) = \frac{1}{2} \left\| \mathbf{Y} - \mu \cdot \mathbf{1} - \sum_{i=1}^p \mathbf{X}_i \theta_i + \sum_{i < j} \mathbf{X}_{i:j} \theta_{i:j} \right\|_2^2, \quad (8)$$

and “logistic loss” (negative Bernoulli log-likelihood) for the binomial response model:

$$\mathcal{L}(\mathbf{Y}; \mu, \theta) = - \left[ \mathbf{Y}^T (\mu \cdot \mathbf{1} + \sum_{i=1}^p \mathbf{X}_i \theta_i + \sum_{i < j} \mathbf{X}_{i:j} \theta_{i:j}) - \mathbf{1}^T \log \left( \mathbf{1} + \exp(\mu \cdot \mathbf{1} + \sum_{i=1}^p \mathbf{X}_i \theta_i + \sum_{i < j} \mathbf{X}_{i:j} \theta_{i:j}) \right) \right], \quad (9)$$

where the log and exp are taken component-wise. Notice that the optimization problem (7) does not impose any hierarchical constraints on the coefficients  $\theta$ , so that its solutions may obey neither weak nor strong hierarchy. This need not be a problem, although it is our contention that some kind of hierarchy is more natural.

## 2.4 Group-lasso and overlapped group-lasso

Since `glinternet`’s workhorse is the group-lasso, we briefly introduce it here. We refer the reader to [Yuan and Lin, 2006] for more technical details.

The group-lasso can be thought of as a more general version of the well-known lasso. Suppose there are  $p$  groups of variables (possibly of different sizes), and let the feature matrix for group  $i$  be denoted by  $\mathbf{X}_i$ . Let  $\mathbf{Y}$  denote the vector of responses. For squared-error loss, the group-lasso estimates  $\{\hat{\beta}_j\}_1^p$  by solving

$$\operatorname{argmin}_{\mu, \beta} \frac{1}{2} \|\mathbf{Y} - \mu \cdot \mathbf{1} - \sum_{j=1}^p \mathbf{X}_j \beta_j\|_2^2 + \lambda \sum_{j=1}^p \gamma_j \|\beta_j\|_2. \quad (10)$$

Note that if each group consists of only one variable, this reduces to the lasso criterion. In our application, each indicator matrix  $\mathbf{X}_j$  will represent a group. If an estimate  $\hat{\beta}_j$  is nonzero, then *all* its components are typically nonzero. Hence the group-lasso applied to the  $\mathbf{X}_j$ ’s intuitively selects those variables that have a strong overall contribution from all their levels toward explaining the response.

The parameter  $\lambda$  controls the amount of regularization, with larger values implying more regularization. The  $\gamma$ ’s allow each group to be penalized to different extents; we set them all equal to 1 (see Section 8). To solve (10), we start with  $\lambda$  just large enough that all estimates are zero. Decreasing  $\lambda$  along a grid of values results in a path of solutions, from which an optimal  $\lambda$  can be chosen by cross validation or some other model selection procedure.

The Karush-Kuhn-Tucker (KKT) optimality conditions for the group-lasso are simple to compute and check. For group  $i$ , they are

$$\|\mathbf{X}_i^T(\mathbf{Y} - \hat{\mathbf{Y}})\|_2 \leq \gamma_i \lambda \quad \text{if } \hat{\beta}_i = 0 \quad (11)$$

$$\|\mathbf{X}_i^T(\mathbf{Y} - \hat{\mathbf{Y}})\|_2 = \gamma_i \lambda \quad \text{if } \hat{\beta}_i \neq 0. \quad (12)$$

The group-lasso is commonly fit by iterative methods — either block coordinate descent, or some form of gradient descent, and convergence can be confirmed by checking the KKT conditions. Details of the algorithm we use can be found in Section 8.

The overlap group-lasso [Jacob et al., 2009] is a variant of the group-lasso where the groups of variables are allowed to have overlaps, i.e. some variables can show up in more than one group. However, each time a variable shows up in a group, it gets a new coefficient. For example, if a variable is included in 3 groups, then it has 3 coefficients that need to be estimated.

### 3 Methodology and results

We want to fit the first order interaction model in a way that obeys strong hierarchy. We show in Section 3.1 how this can be achieved by adding an overlapped group-lasso penalty to the objective in (7), in addition to the sum-to-zero constraints on the  $\theta$ s. We then show how this optimization problem can be conveniently solved via a group-lasso without overlaps, and without the sum-to-zero constraints.

#### 3.1 Strong hierarchy through overlapped group-lasso

Adding an overlapped group-lasso penalty to (7) is one way of obtaining solutions that satisfy the strong hierarchy property. The results that follow hold for both squared error and logistic loss, but we focus on the former for clarity.

Consider the case where there are two categorical variables  $F_1$  and  $F_2$  with  $L_1$  and  $L_2$  levels respectively. Their indicator matrices are given by  $\mathbf{X}_1$  and  $\mathbf{X}_2$ . (These results generalize trivially to the case of pairwise interactions



with more than 2 variables). We solve

$$\begin{aligned} \operatorname{argmin}_{\mu, \alpha, \tilde{\alpha}} \frac{1}{2} & \left\| \mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{X}_1 \alpha_1 - \mathbf{X}_2 \alpha_2 - [\mathbf{X}_1 \ \mathbf{X}_2 \ \mathbf{X}_{1:2}] \begin{bmatrix} \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \\ \alpha_{1:2} \end{bmatrix} \right\|_2^2 \\ & + \lambda \left( \|\alpha_1\|_2 + \|\alpha_2\|_2 + \sqrt{L_2 \|\tilde{\alpha}_1\|_2^2 + L_1 \|\tilde{\alpha}_2\|_2^2 + \|\alpha_{1:2}\|_2^2} \right) \end{aligned} \quad (13)$$

subject to

$$\sum_{i=1}^{L_1} \alpha_1^i = 0, \quad \sum_{j=1}^{L_2} \alpha_2^j = 0, \quad \sum_{i=1}^{L_1} \tilde{\alpha}_1^i = 0, \quad \sum_{j=1}^{L_2} \tilde{\alpha}_2^j = 0 \quad (14)$$

and

$$\sum_{i=1}^{L_1} \alpha_{1:2}^{ij} = 0 \text{ for fixed } j, \quad \sum_{j=1}^{L_2} \alpha_{1:2}^{ij} = 0 \text{ for fixed } i. \quad (15)$$

Notice that  $\mathbf{X}_i$ ,  $i = 1, 2$  each have two different coefficient vectors  $\alpha_i$  and  $\tilde{\alpha}_i$ , resulting in an overlapped penalty. The  $\sqrt{L_2 \|\tilde{\alpha}_1\|_2^2 + L_1 \|\tilde{\alpha}_2\|_2^2 + \|\alpha_{1:2}\|_2^2}$  term results in estimates that satisfy strong hierarchy, because either  $\hat{\tilde{\alpha}}_1 = \hat{\tilde{\alpha}}_2 = \hat{\alpha}_{1:2} = 0$  or all are nonzero, i.e. interactions are always present with both main effects.

The constants  $L_1$  and  $L_2$  are chosen to put  $\tilde{\alpha}_1$ ,  $\tilde{\alpha}_2$ , and  $\alpha_{1:2}$  on the same scale. To motivate this, note that we can write

$$X_1 \tilde{\alpha}_1 = X_{1:2} \underbrace{[\tilde{\alpha}_1, \dots, \tilde{\alpha}_1]^T}_{L_2 \text{ copies}}, \quad (16)$$

and similarly for  $X_2 \tilde{\alpha}_2$ . We now have a representation for  $\tilde{\alpha}_1$  and  $\tilde{\alpha}_2$  with respect to the space defined by  $X_{1:2}$ , so that they are “comparable” to  $\alpha_{1:2}$ . We then have

$$\| \underbrace{[\tilde{\alpha}_1, \dots, \tilde{\alpha}_1]}_{L_2 \text{ copies}} \|_2^2 = L_2 \|\tilde{\alpha}_1\|_2^2 \quad (17)$$

and likewise for  $\tilde{\alpha}_2$ . More details are given in Section 3.2 below.

The actual main effects and interactions can be recovered as

$$\hat{\theta}_1 = \hat{\alpha}_1 + \hat{\tilde{\alpha}}_1 \quad (18)$$

$$\hat{\theta}_2 = \hat{\alpha}_2 + \hat{\tilde{\alpha}}_2 \quad (19)$$

$$\hat{\theta}_{1:2} = \hat{\alpha}_{1:2}. \quad (20)$$

Because of the strong hierarchy property mentioned above, we also have

$$\hat{\theta}_{1:2} \neq 0 \implies \hat{\theta}_1 \neq 0 \text{ and } \hat{\theta}_2 \neq 0. \quad (21)$$

In standard anova, there are a number of ways of dealing with the over-parametrization for factors, other than the sum-to-zero constraints we use here; for example, setting the coefficient for the first or last level to zero, or other chosen contrasts. These also lead to a reduced parameterization. While this choice does not alter the fit in the standard case, with parameter regularization it does. Our choice is symmetric and appears to be natural, in that it treats all the coefficients equally.

### 3.2 Equivalence with unconstrained group-lasso

We now show how to solve the constrained overlapped group-lasso problem by solving an equivalent *unconstrained* group-lasso problem — a simplification with considerable computational advantages.

We will need two lemmas. The first shows that because we fit an intercept in the model, the estimated coefficients  $\hat{\beta}$  for any categorical variables in the model will have mean zero.

**Lemma 1.** *Let  $X$  be an indicator matrix. Then the solution  $\hat{\beta}$  to*

$$\operatorname{argmin}_{\mu, \beta} \frac{1}{2} \|\mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_2 \quad (22)$$

*satisfies*

$$\bar{\hat{\beta}} = 0. \quad (23)$$

*The same is true for logistic loss.*

*Proof.* Because  $\mathbf{X}$  is an indicator matrix, each row consists of exactly a single 1 (all other entries 0), so that

$$\mathbf{X} \cdot c\mathbf{1} = c\mathbf{1} \quad (24)$$

for any constant  $c$ . It follows that if  $\hat{\mu}$  and  $\hat{\beta}$  are solutions, then so are  $\hat{\mu} + c\mathbf{1}$  and  $\hat{\beta} - c\mathbf{1}$ . But the norm  $\|\hat{\beta} - c\mathbf{1}\|_2$  is minimized for  $c = \bar{\hat{\beta}}$ .  $\square$

Note that this would be true for each of any number of  $\mathbf{X}_j$  and  $\beta_j$  in the model, as long as there is an unpenalized intercept.

The next Lemma states that if we include two intercepts in the model, one penalized and the other unpenalized, then the penalized intercept will be estimated to be zero. This is because we can achieve the same fit with a lower penalty by taking  $\mu \longleftarrow \mu + \tilde{\mu}$ .

**Lemma 2.** *The optimization problem*

$$\operatorname{argmin}_{\mu, \tilde{\mu}, \beta} \frac{1}{2} \|\mathbf{Y} - \mu \cdot \mathbf{1} - \tilde{\mu} \cdot \mathbf{1} - \mathbf{X}\beta\|_2^2 + \lambda \sqrt{\tilde{\mu}^2 + \|\beta\|_2^2} \quad (25)$$

has solution  $\hat{\mu} = 0$  for all  $\lambda > 0$ . The same result holds for logistic loss.

The next theorem shows how the overlapped group-lasso in Section 3.1 reduces to a group-lasso.

**Theorem 1.** *Solving the constrained optimization problem (13) - (15) in Section 3.1 is equivalent to solving the unconstrained problem*

$$\operatorname{argmin}_{\mu, \beta} \frac{1}{2} \|\mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{X}_1\beta_1 - \mathbf{X}_2\beta_2 - \mathbf{X}_{1:2}\beta_{1:2}\|_2^2 + \lambda (\|\beta_1\|_2 + \|\beta_2\|_2 + \|\beta_{1:2}\|_2). \quad (26)$$

*Proof.* We need to show that the group-lasso objective can be equivalently written as an overlapped group-lasso with the appropriate constraints on the parameters. We begin by rewriting (13) as

$$\operatorname{argmin}_{\mu, \tilde{\mu}, \alpha, \tilde{\alpha}} \frac{1}{2} \left\| \mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{X}_1\alpha_1 - \mathbf{X}_2\alpha_2 - [\mathbf{1} \ \mathbf{X}_1 \ \mathbf{X}_2 \ \mathbf{X}_{1:2}] \begin{bmatrix} \tilde{\mu} \\ \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \\ \alpha_{1:2} \end{bmatrix} \right\|_2^2 + \lambda \left( \|\alpha_1\|_2 + \|\alpha_2\|_2 + \sqrt{L_1 L_2 \tilde{\mu}^2 + L_2 \|\tilde{\alpha}_1\|_2^2 + L_1 \|\tilde{\alpha}_2\|_2^2 + \|\alpha_{1:2}\|_2^2} \right). \quad (27)$$

By Lemma 2, we will estimate  $\hat{\mu} = 0$ . Therefore we have not changed the solutions in any way.

Lemma 1 shows that the first two constraints in (14) are satisfied by the estimated main effects  $\hat{\beta}_1$  and  $\hat{\beta}_2$ . We now show that

$$\|\beta_{1:2}\|_2 = \sqrt{L_1 L_2 \tilde{\mu}^2 + L_2 \|\tilde{\alpha}_1\|_2^2 + L_1 \|\tilde{\alpha}_2\|_2^2 + \|\alpha_{1:2}\|_2^2} \quad (28)$$

where the  $\tilde{\alpha}_1$ ,  $\tilde{\alpha}_2$ , and  $\alpha_{1:2}$  satisfy the constraints in (14) and (15), and along with  $\tilde{\mu}$  represent a reparameterization of  $\beta_{1:2}$ .

For fixed levels  $i$  and  $j$ , we can decompose  $\beta_{1:2}$  (see [Scheffe, 1959]) as

$$\begin{aligned} \beta_{1:2}^{ij} &= \beta_{1:2}^{\cdot\cdot} + (\beta_{1:2}^{i\cdot} - \beta_{1:2}^{\cdot\cdot}) + (\beta_{1:2}^{\cdot j} - \beta_{1:2}^{\cdot\cdot}) + (\beta_{1:2}^{ij} - \beta_{1:2}^{i\cdot} - \beta_{1:2}^{\cdot j} + \beta_{1:2}^{\cdot\cdot}) \\ &\equiv \tilde{\mu} + \tilde{\alpha}_1^i + \tilde{\alpha}_2^j + \alpha_{1:2}^{ij}. \end{aligned}$$

It follows that the whole  $(L_1 L_2)$ -vector  $\beta_{1:2}$  can be written as

$$\beta_{1:2} = \mathbf{1}\tilde{\mu} + \mathbf{Z}_1\tilde{\alpha}_1 + \mathbf{Z}_2\tilde{\alpha}_2 + \alpha_{1:2}, \quad (29)$$

where  $\mathbf{Z}_1$  is a  $L_1 L_2 \times L_1$  indicator matrix of the form

$$\underbrace{\begin{pmatrix} \mathbf{1}_{L_2 \times 1} & 0 & \cdots & 0 \\ 0 & \mathbf{1}_{L_2 \times 1} & \cdots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \mathbf{1}_{L_2 \times 1} \end{pmatrix}}_{L_1 \text{ columns}} \quad (30)$$

and  $\mathbf{Z}_2$  is a  $L_1 L_2 \times L_2$  indicator matrix of the form

$$L_1 \text{ copies} \left\{ \begin{pmatrix} I_{L_2 \times L_2} \\ \vdots \\ I_{L_2 \times L_2} \end{pmatrix} \right\} \quad (31)$$

It follows that

$$\mathbf{Z}_1\tilde{\alpha}_1 = \underbrace{(\tilde{\alpha}_1^1, \dots, \tilde{\alpha}_1^1)}_{L_2 \text{ copies}} \underbrace{(\tilde{\alpha}_1^2, \dots, \tilde{\alpha}_1^2)}_{L_2 \text{ copies}}, \dots, \underbrace{(\tilde{\alpha}_1^{L_1}, \dots, \tilde{\alpha}_1^{L_1})}_{L_2 \text{ copies}}^T \quad (32)$$

and

$$\mathbf{Z}_2\tilde{\alpha}_2 = (\tilde{\alpha}_2^1, \dots, \tilde{\alpha}_2^{L_2}, \tilde{\alpha}_2^1, \dots, \tilde{\alpha}_2^{L_2}, \dots, \tilde{\alpha}_2^1, \dots, \tilde{\alpha}_2^{L_2})^T. \quad (33)$$

Note that  $\tilde{\alpha}_1$ ,  $\tilde{\alpha}_2$ , and  $\alpha_{1:2}$ , by definition, satisfy the constraints (14) and (15). This can be used to show, by direct calculation, that the four additive components in (29) are mutually orthogonal, so that we can write

$$\|\beta_{1:2}\|_2^2 = \|\mathbf{1}\tilde{\mu}\|_2^2 + \|\mathbf{Z}_1\tilde{\alpha}_1\|_2^2 + \|\mathbf{Z}_2\tilde{\alpha}_2\|_2^2 + \|\alpha_{1:2}\|_2^2 \quad (34)$$

$$= L_1 L_2 \tilde{\mu}^2 + L_2 \|\tilde{\alpha}_1\|_2^2 + L_1 \|\tilde{\alpha}_2\|_2^2 + \|\alpha_{1:2}\|_2^2. \quad (35)$$

We have shown that the penalty in the group-lasso problem is equivalent to the penalty in the constrained overlapped group-lasso. It remains to show that the loss functions in both problems are also the same. Since  $\mathbf{X}_{1:2}\mathbf{Z}_1 = \mathbf{X}_1$  and  $\mathbf{X}_{1:2}\mathbf{Z}_2 = \mathbf{X}_2$ , this can be seen by a direct computation:

$$\mathbf{X}_{1:2}\beta_{1:2} = \mathbf{X}_{1:2}(\mathbf{1}\tilde{\mu} + \mathbf{Z}_1\tilde{\alpha}_1 + \mathbf{Z}_2\tilde{\alpha}_2 + \alpha_{1:2}) \quad (36)$$

$$= \mathbf{1}\tilde{\mu} + \mathbf{X}_1\tilde{\alpha}_1 + \mathbf{X}_2\tilde{\alpha}_2 + \mathbf{X}_{1:2}\alpha_{1:2} \quad (37)$$

$$= [\mathbf{1} \ \mathbf{X}_1 \ \mathbf{X}_2 \ \mathbf{X}_{1:2}] \begin{bmatrix} \tilde{\mu} \\ \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \\ \alpha_{1:2} \end{bmatrix} \quad (38)$$

□

Theorem 1 shows that we can use the group-lasso to obtain estimates that satisfy strong hierarchy, without solving the overlapped group-lasso with constraints. The theorem also shows that the main effects and interactions can be extracted with

$$\hat{\theta}_1 = \hat{\beta}_1 + \hat{\alpha}_1 \quad (39)$$

$$\hat{\theta}_2 = \hat{\beta}_2 + \hat{\alpha}_2 \quad (40)$$

$$\hat{\theta}_{1:2} = \hat{\alpha}_{1:2}. \quad (41)$$

Although this theorem is stated and proved with just two categorical variables  $F_1$  and  $F_2$ , the same results extends in a trivial way to all pairs of interactions with  $p$  variables. The only small modification is that a main effect will potentially get contributions from more than one interaction pair.

We discuss the properties of the `glinternet` estimates in the next section.

### 3.3 Properties of the `glinternet` penalty

While `glinternet` treats the problem as a group-lasso, examining the equivalent overlapped group-lasso version makes it easier to draw insights about the behaviour of the method under various scenarios. Recall that the overlapped penalty for two variables is given by

$$\|\alpha_1\|_2 + \|\alpha_2\|_2 + \sqrt{L_2\|\tilde{\alpha}_1\|_2^2 + L_1\|\tilde{\alpha}_2\|_2^2 + \|\alpha_{1:2}\|_2^2}. \quad (42)$$

We consider a noiseless scenario. If the ground truth is additive, i.e.  $\alpha_{1:2} = 0$ , then  $\tilde{\alpha}_1$  and  $\tilde{\alpha}_2$  will be estimated to be zero. This is because for  $L_1, L_2 \geq 2$  and  $a, b \geq 0$ , we have

$$\sqrt{L_2a^2 + L_1b^2} \geq a + b. \quad (43)$$

Thus it is advantageous to place all the main effects in  $\alpha_1$  and  $\alpha_2$ , because doing so results in a smaller penalty. Therefore, if the truth has no interactions, then `glinternet` picks out only main effects.

If an interaction was present ( $\alpha_{1:2} > 0$ ), the derivative of the penalty term with respect to  $\alpha_{1:2}$  is

$$\frac{\alpha_{1:2}}{\sqrt{L_2\|\tilde{\alpha}_1\|_2^2 + L_1\|\tilde{\alpha}_2\|_2^2 + \|\alpha_{1:2}\|_2^2}}. \quad (44)$$

The presence of main effects allows this derivative to be smaller, thus allowing the algorithm to pay a smaller penalty (as compared to no main effects present) for making  $\hat{\alpha}_{1:2}$  nonzero. This shows interactions whose main effects are also present are discovered before pure interactions.

### 3.4 Interaction between a categorical variable and a continuous variable

We describe how to extend Theorem 1 to interaction between a continuous variable and a categorical variable.

Consider the case where we have a categorical variable  $F$  with  $L$  levels, and a continuous variable  $Z$ . Let  $\mu_i = \mathbb{E}[Y|F = i, Z = z]$ . There are four cases:

- $\mu_i = \mu$  (no main effects, no interactions)
- $\mu_i = \mu + \theta_1^i$  (main effect  $F$ )
- $\mu_i = \mu + \theta_1^i + \theta_2 z$  (two main effects)
- $\mu_i = \mu + \theta_1^i + \theta_2 z + \theta_{1:2}^i z$  (main effects and interaction)

As before, we impose the constraints  $\sum_{i=1}^L \theta_1^i = 0$  and  $\sum_{i=1}^L \theta_{1:2}^i = 0$ . An overlapped group-lasso of the form

$$\begin{aligned} \operatorname{argmin}_{\mu, \alpha, \tilde{\alpha}} \frac{1}{2} \left\| \mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{X}\alpha_1 - \mathbf{Z}\alpha_2 - [\mathbf{X} \quad \mathbf{Z} \quad (\mathbf{X} * \mathbf{Z})] \begin{bmatrix} \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \\ \alpha_{1:2} \end{bmatrix} \right\|_2^2 \\ + \lambda \left( \|\alpha_1\|_2 + |\alpha_2| + \sqrt{\|\tilde{\alpha}_1\|_2^2 + L\tilde{\alpha}_2^2 + \|\alpha_{1:2}\|_2^2} \right) \end{aligned} \quad (45)$$

subject to

$$\sum_{i=1}^L \alpha_1^i = 0, \quad \sum_{i=1}^L \tilde{\alpha}_1^i = 0, \quad \sum_{i=1}^L \alpha_{1:2}^i = 0 \quad (46)$$

allows us to obtain estimates of the interaction term that satisfy strong hierarchy. This is again due to the nature of the square root term in the penalty. The actual main effects and interactions can be recovered as

$$\hat{\theta}_1 = \hat{\alpha}_1 + \hat{\tilde{\alpha}}_1 \quad (47)$$

$$\hat{\theta}_2 = \hat{\alpha}_2 + \hat{\tilde{\alpha}}_2 \quad (48)$$

$$\hat{\theta}_{1:2} = \hat{\alpha}_{1:2}. \quad (49)$$

We have the following extension of Theorem 1:

**Theorem 2.** *Solving the constrained overlapped group-lasso above is equivalent to solving*

$$\begin{aligned} \operatorname{argmin}_{\mu, \beta} \frac{1}{2} \|\mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{X}\beta_1 - \mathbf{Z}\beta_2 - (\mathbf{X} * [\mathbf{1} \quad \mathbf{Z}])\beta_{1:2}\|_2^2 \\ + \lambda (\|\beta_1\|_2 + |\beta_2| + \|\beta_{1:2}\|_2). \end{aligned} \quad (50)$$

*Proof.* We proceed as in the proof of Theorem 1 and introduce an additional parameter  $\tilde{\mu}$  into the overlapped objective:

$$\begin{aligned} \operatorname{argmin}_{\mu, \tilde{\mu}, \alpha, \tilde{\alpha}} \frac{1}{2} \left\| \mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{X}\alpha_1 - \mathbf{Z}\alpha_2 - [\mathbf{1} \quad \mathbf{X} \quad \mathbf{Z} \quad (\mathbf{X} * \mathbf{Z})] \begin{bmatrix} \tilde{\mu} \\ \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \\ \alpha_{1:2} \end{bmatrix} \right\|_2^2 \\ + \lambda \left( \|\alpha_1\|_2 + |\alpha_2| + \sqrt{L\tilde{\mu}^2 + \|\tilde{\alpha}_1\|_2^2 + L\tilde{\alpha}_2^2 + \|\alpha_{1:2}\|_2^2} \right) \end{aligned} \quad (51)$$

As before, this does not change the solutions because we will have  $\hat{\mu} = 0$  (see Lemma 2).

Decompose the  $2L$ -vector  $\beta_{1:2}$  into

$$\begin{bmatrix} \eta_1 \\ \eta_2 \end{bmatrix}, \quad (52)$$

where  $\eta_1$  and  $\eta_2$  both have dimension  $L \times 1$ . Apply the anova decomposition to both to obtain

$$\eta_1^i = \eta_1 + (\eta_1^i - \eta_1) \quad (53)$$

$$\equiv \tilde{\mu} + \tilde{\alpha}_1^i \quad (54)$$

and

$$\eta_2^i = \eta_2 + (\eta_2^i - \eta_2) \quad (55)$$

$$\equiv \tilde{\alpha}_2 + \alpha_{1:2}^i. \quad (56)$$

Note that  $\tilde{\alpha}_1$  is a  $(L \times 1)$ -vector that satisfies  $\sum_{i=1}^L \tilde{\alpha}_1^i = 0$ , and likewise for  $\alpha_{1:2}$ . This allows us to write

$$\beta_{1:2} = \begin{bmatrix} \tilde{\mu} \cdot \mathbf{1}_{L \times 1} \\ \tilde{\alpha}_2 \cdot \mathbf{1}_{L \times 1} \end{bmatrix} + \begin{bmatrix} \tilde{\alpha}_1 \\ \alpha_{1:2} \end{bmatrix} \quad (57)$$

It follows that

$$\|\beta_{1:2}\|_2^2 = L\tilde{\mu}^2 + \|\tilde{\alpha}_1\|_2^2 + L\tilde{\alpha}_2^2 + \|\alpha_{1:2}\|_2^2, \quad (58)$$

which shows that the penalties in both problems are equivalent. A direct computation shows that the loss functions are also equivalent:

$$(\mathbf{X} * [\mathbf{1} \quad \mathbf{Z}])\beta_{1:2} = [\mathbf{X} \quad (\mathbf{X} * \mathbf{Z})]\beta_{1:2} \quad (59)$$

$$= [\mathbf{X} \quad (\mathbf{X} * \mathbf{Z})] \left( \begin{bmatrix} \tilde{\mu} \cdot \mathbf{1}_{L \times 1} \\ \tilde{\alpha}_2 \cdot \mathbf{1}_{L \times 1} \end{bmatrix} + \begin{bmatrix} \tilde{\alpha}_1 \\ \alpha_{1:2} \end{bmatrix} \right) \quad (60)$$

$$= \tilde{\mu} \cdot \mathbf{1} + \mathbf{X}\tilde{\alpha}_1 + \mathbf{Z}\tilde{\alpha}_2 + (\mathbf{X} * \mathbf{Z})\alpha_{1:2}. \quad (61)$$

□

Theorem 2 allows us to accommodate interactions between continuous and categorical variables by simply parametrizing the interaction term as  $\mathbf{X} * [\mathbf{1} \quad \mathbf{Z}]$ , where  $\mathbf{X}$  is the indicator matrix representation for categorical variables that we have been using all along. We then proceed as before with a group-lasso.

### 3.5 Interaction between two continuous variables

We have seen that the appropriate representations for the interaction terms are

- $\mathbf{X}_1 * \mathbf{X}_2 = \mathbf{X}_{1:2}$  for categorical variables
- $\mathbf{X} * [\mathbf{1} \quad \mathbf{Z}] = [\mathbf{X} \quad (\mathbf{X} * \mathbf{Z})]$  for one categorical variable and one continuous variable.

How should we represent the interaction between two continuous variables? We follow the traditional approach, and use the bilinear interaction term  $Z_1 \cdot Z_2$ . We can use exactly the same framework as before:

$$\mathbf{Z}_{1:2} = [\mathbf{1} \quad \mathbf{Z}_1] * [\mathbf{1} \quad \mathbf{Z}_2] \quad (62)$$

$$= [\mathbf{1} \quad \mathbf{Z}_1 \quad \mathbf{Z}_2 \quad (\mathbf{Z}_1 * \mathbf{Z}_2)]. \quad (63)$$

This is indeed the case. A linear interaction model for  $\mathbf{Z}_1$  and  $\mathbf{Z}_2$  is given by

$$\mathbb{E}[Y|Z_1 = z_1, Z_2 = z_2] = \mu + \theta_1 z_1 + \theta_2 z_2 + \theta_{1:2} z_1 z_2. \quad (64)$$

Unlike the previous cases where there were categorical variables, there are no constraints on any of the coefficients. It follows that the overlapped



group-lasso

$$\operatorname{argmin}_{\mu, \tilde{\mu}, \alpha, \tilde{\alpha}} \frac{1}{2} \left\| \mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{Z}_1 \alpha_1 - \mathbf{Z}_2 \alpha_2 - [\mathbf{1} \quad \mathbf{Z}_1 \quad \mathbf{Z}_2 \quad (\mathbf{Z}_1 * \mathbf{Z}_2)] \begin{bmatrix} \tilde{\mu} \\ \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \\ \alpha_{1:2} \end{bmatrix} \right\|_2^2 + \lambda \left( |\alpha_1| + |\alpha_2| + \sqrt{\tilde{\mu}^2 + \tilde{\alpha}_1^2 + \tilde{\alpha}_2^2 + \alpha_{1:2}^2} \right) \quad (65)$$

is trivially equivalent to

$$\operatorname{argmin}_{\mu, \beta} \frac{1}{2} \left\| \mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{Z}_1 \beta_1 - \mathbf{Z}_2 \beta_2 - ([\mathbf{1} \quad \mathbf{Z}_1] * [\mathbf{1} \quad \mathbf{Z}_2]) \beta_{1:2} \right\|_2^2 + \lambda (|\beta_1| + |\beta_2| + \|\beta_{1:2}\|_2), \quad (66)$$

with the  $\beta$ 's taking the place of the  $\alpha$ 's. Note that we will have  $\hat{\mu} = 0$ .

## 4 Variable screening


`glinternet` works by solving a group-lasso with  $p + \binom{p}{2}$  groups of variables.

Even for moderate  $p$  ( $\sim 10^5$ ), we will require some form of screening to reduce the dimension of the interaction search space. We have argued that models satisfying hierarchy make sense, so that it is natural to consider screening devices that hedge on the presence of main effects. We discuss two screening methods in this section: gradient boosting, and an adaptive screen based on the strong rules of [Tibshirani et al, 2012]. We describe the boosting approach first.

### 4.1 Screening with boosted trees

AdaBoost [Freund and Schapire, 1995] and gradient boosting [Friedman, 2000] are effective approaches for building ensembles of weak learners such as decision trees. One of the advantages of trees is that they are able to model nonlinear effects and high-order interactions. For example, a depth-2 tree essentially represents an interaction between the variables involved in the two splits, which suggests that boosting with depth-2 trees is a way of building a first-order interaction model. Note that the interactions are hierarchical, because in finding the optimal first split, the boosting algorithm is looking for the best main effect. The subsequent split is then made, conditioned on the first split.

If we boost with  $T$  trees, then we end up with a model that has at most  $T$  interaction pairs. The following diagram gives a schematic of the boosting iterations with categorical variables.



In the first tree, levels 2 and 3 of  $F_1$  are not involved in the interaction with  $F_2$ . Therefore each tree in the boosted model does not represent an interaction among all the levels of the two variables, but only among a subset of the levels. To enforce the full interaction structure, one could use fully-split trees, but we do not develop this approach for two reasons. First, boosting is a sequential procedure and is quite slow even for moderately sized problems. Using fully split trees will further degrade its runtime. Second, in variables with many levels, it is reasonable to expect that the interactions only occur among a few of the levels. If this were true, then a complete interaction that is weak for every combination of levels might be selected over a strong partial interaction. But it is the strong partial interaction that we are interested in.

Boosting is feasible because it is a greedy algorithm. If  $p$  is the number of variables, an exhaustive search involves  $\mathcal{O}(p^2)$  variables, whereas boosting operates with  $\mathcal{O}(p)$ . To use the boosted model as a screening device for interaction candidates, we take the set of all unique interactions from the collection of trees. For example, in our schematic above, we would add  $F_{1:2}$  and  $F_{11:23}$  to our candidate set of interactions.

In our experiments, using boosting as a screen did not perform as well as we hoped. There is the issue of selecting tuning parameters such as the amount of shrinkage and the number of trees to use. Lowering the shrinkage and increasing the number of trees improves false discovery rates, but at a significant cost to speed. In the next section, we describe a screening approach that is based on computing inner products that is efficient and that can be integrated with the strong rules for the group lasso.

## 4.2 An adaptive screening procedure

The strong rules [Tibshirani et al, 2012] for lasso-type problems are effective heuristics for discarding large numbers of variables that are likely to be redundant. As a result, the strong rules can dramatically speed up the convergence of algorithms because they can concentrate on a smaller set (we call this the *strong set*) of variables that are more likely to be nonzero. The strong rules are not safe, however, meaning that it is possible that some of the discarded variables are actually supposed to be nonzero. Because of this, after our algorithm has converged on the strong set, we have to check the KKT conditions on the discarded set. Those variables that do not satisfy the conditions then have to be added to the current set of nonzero variables, and we fit on this expanded set. This happens rarely in our experience, i.e. the discarded variables tend to remain zero after the algorithm has converged on the strong set, which means we rarely have to do multiple rounds of fitting for any given value of the regularization parameter  $\lambda$ .

Suppose we have a solution  $\hat{\mathbf{Y}}_\ell$  at  $\lambda_\ell$ , and wish to compute the solution at  $\lambda_{\ell+1} < \lambda_\ell$ . The strong rule for the group-lasso involves computing  $s_i = \|\mathbf{X}_i^T(\mathbf{Y} - \hat{\mathbf{Y}}_\ell)\|_2$  for every group of variables  $\mathbf{X}_i$ . We know from the KKT optimality conditions (see Section 2.4) that

$$\begin{aligned} s_i &\leq \lambda_\ell & \text{if } \hat{\beta}_i &= 0 \\ s_i &= \lambda_\ell & \text{if } \hat{\beta}_i &\neq 0, \end{aligned} \tag{68}$$

(assuming the  $\gamma_i$  in (11) are all 1). When we decrease  $\lambda$  to  $\lambda_{\ell+1}$ , we expect to increase the active set. The strong rule discards a group  $i$  if  $s_i < \lambda_{\ell+1} - (\lambda_\ell - \lambda_{\ell+1})$ . If this is computationally feasible for all  $p + \binom{p}{2}$  groups, then there is no need for screening; we simply fit the group-lasso on those groups that pass the strong rules filter. Otherwise, we approximate this by screening only on the groups that correspond to main effects. We then take the candidate set of interactions to consist of all pairwise interactions between the variables that passed this screen. Note that we will have already computed the  $s_i$  for the strong rules from checking the KKT conditions for the solutions at  $\lambda_\ell$ . This allows us to integrate screening with the strong rules in an efficient manner.

An example will illustrate.

Suppose we have 10,000 variables ( $\sim 50 \times 10^6$  possible interactions), but we are computationally limited to a group-lasso with  $10^6$  groups. Assume we have the fit for  $\lambda = \lambda_\ell$ , and want to move on to  $\lambda_{\ell+1}$ . Let  $r_{\lambda_\ell} = \mathbf{Y} - \hat{\mathbf{Y}}_{\lambda_\ell}$  denote the current residual. At this point, the variable scores  $s_i =$

$\|\mathbf{X}_i^T r_{\lambda_\ell}\|_2$  have already been computed from checking the KKT conditions at the solutions for  $\lambda_\ell$ . We restrict ourselves to the 10,000 variables, and take the 100 with the highest scores. Denote this set by  $\mathcal{T}_{100}^{\lambda_{\ell+1}}$ . The candidate set of variables for the group-lasso is then given by  $\mathcal{T}_{100}^{\lambda_{\ell+1}}$  together with the pairwise interactions between *all* 10,000 variables and  $\mathcal{T}_{100}^{\lambda_{\ell+1}}$ . Because this gives a candidate set with about  $100 \times 10,000 = 10^6$  terms, the computation is now feasible. We then compute the group-lasso on this candidate set, and repeat the procedure with the new residual  $r_{\lambda_{\ell+1}}$ .

This screen is easy to compute since it is based on inner products. Moreover, they can be computed in parallel. The procedure also integrates well with the strong rules by reusing inner products computed from the fit for a previous  $\lambda$ .

## 5 Related work and approaches

We describe some past and related approaches to discovering interactions. We give a short synopsis of how they work, and say why they are inadequate for our purposes. The method most similar to ours is **hierNet**.

### 5.1 Logic regression [Ruczinski I, 2003]

Logic regression finds boolean combinations of variables that have high predictive power of the response variable. For example, a combination might look like

$$(F_1 \text{ and } F_3) \text{ or } F_5. \tag{69}$$

This is an example of an interaction that is of higher-order than what **glinternet** handles, and is an appealing aspect of logic regression. However, logic regression does not accommodate continuous variables or categorical variables with more than two levels. We do not make comparisons with logic regression in our simulations for this reason.

### 5.2 Composite absolute penalties [Zhao et al., 2009]

Like **glinternet**, this is also a penalty-based approach. CAP employs penalties of the form

$$\|(\beta_i, \beta_j)\|_{\gamma_1} + \|\beta_j\|_{\gamma_2} \tag{70}$$

where  $\gamma_1 > 1$ . Such a penalty ensures that  $\hat{\beta}_i \neq 0$  whenever  $\hat{\beta}_j \neq 0$ . It is possible that  $\hat{\beta}_i \neq 0$  but  $\hat{\beta}_j = 0$ . In other words, the penalty makes  $\hat{\beta}_j$  hierarchically dependent on  $\hat{\beta}_i$ : it can only be nonzero after  $\hat{\beta}_i$  becomes nonzero. It is thus possible to use CAP penalties to build interaction models that satisfy hierarchy. For example, a penalty of the form  $\|(\theta_1, \theta_2, \theta_{1:2})\|_2 + \|\theta_{1:2}\|_2$  will result in estimates that satisfy  $\hat{\theta}_{1:2} \neq 0 \implies \hat{\theta}_1 \neq 0$  and  $\hat{\theta}_2 \neq 0$ . We can thus build a linear interaction model for two categorical variables by solving

$$\begin{aligned} \operatorname{argmin}_{\mu, \theta} \frac{1}{2} \|\mathbf{Y} - \mu \cdot \mathbf{1} - \mathbf{X}_1 \theta_1 - \mathbf{X}_2 \theta_2 - \mathbf{X}_{1:2} \theta_{1:2}\|_2^2 \\ + \lambda (\|(\theta_1, \theta_2, \theta_{1:2})\|_2 + \|\theta_{1:2}\|_2) \end{aligned} \quad (71)$$

subject to (3) and (4). We see that the CAP approach differs from `glinternet` in that we have to solve a constrained optimization problem which is considerably more complicated, thus making it unclear if CAP will be computationally feasible for larger problems. The form of the penalties are also different: the interaction coefficient in CAP is penalized twice, whereas `glinternet` penalizes it once. It is not obvious what the relationship between the two algorithms' solutions would be.

### 5.3 hierNet [Bien et al, 2013]

This is a method that, like `glinternet`, seeks to find interaction estimates that obey hierarchy with regularization. The optimization problem that `hierNet` solves is

$$\operatorname{argmin}_{\mu, \beta, \theta} \frac{1}{2} \sum_{i=1}^n (y_i - \mu - x_i^T \beta - \frac{1}{2} x_i^T \theta x_i)^2 + \lambda \mathbf{1}^T (\beta^+ + \beta^-) + \frac{\lambda}{2} \|\theta\|_1 \quad (72)$$

subject to

$$\theta = \theta^T, \|\theta_j\|_1 \leq \beta_j^+ + \beta_j^-, \beta_j^+ \geq 0, \beta_j^- \geq 0. \quad (73)$$

The main effects are represented by  $\beta$ , and interactions are given by  $\theta$ . The first constraint enforces symmetry in the interaction coefficients.  $\beta_j^+$  and  $\beta_j^-$  are the positive and negative parts of  $\beta_j$ , and are given by  $\beta_j^+ = \max(0, \beta_j)$  and  $\beta_j^- = -\min(0, \beta_j)$  respectively. The constraint  $\|\theta_j\|_1 \leq \beta_j^+ + \beta_j^-$  implies that if some components of the  $j$ -th row of  $\theta$  are estimated to be nonzero, then the main effect  $\beta_j$  will also be estimated to be nonzero. Since  $\theta_j$  corresponds to interactions between the  $j$ -th variable and all the

other variables, this implies that the solutions to the `hierNet` objective satisfy weak hierarchy. One can think of  $\beta_j^+ + \beta_j^-$  as a budget for the amount of interactions that are allowed to be nonzero.

The `hierNet` objective can be modified to obtain solutions that satisfy strong hierarchy, which makes it in principle comparable to `glinternet`. Currently, `hierNet` is only able to accommodate binary and continuous variables, and is practically limited to fitting models with fewer than 1000 variables.

## 6 Simulation study

We perform simulations to see if `glinternet` is competitive with existing methods. `hierNet` is a natural benchmark because it also tries to find interactions subject to hierarchical constraints. Because `hierNet` only works with continuous variables and 2-level categorical variables, we include gradient boosting as a competitor for the scenarios where `hierNet` cannot be used.

### 6.1 False discovery rates

We simulate 4 different setups:

1. Truth obeys strong hierarchy. The interactions are only among pairs of nonzero main effects.
2. Truth obeys weak hierarchy. Each interaction has only one of its main effects present.
3. Truth is anti-hierarchical. The interactions are only among pairs of main effects that are not present.
4. Truth is pure interaction. There are no main effects present, only interactions.

Each case is generated with  $n = 500$  observations and  $p = 30$  continuous variables, with a signal to noise ratio of 1. Where applicable, there are 10 main effects and/or 10 interactions in the ground truth. The interaction and main effect coefficients are sampled from  $N(0, 1)$ , so that the variance in the observations should be split equally between main effects and interactions.

Boosting is done with 5000 depth-2 trees and a learning rate of 0.001. Each tree represents a candidate interaction, and we can compute the improvement to fit due to this candidate pair. Summing up the improvement

over the 5000 trees gives a score for each interaction pair, which can then be used to order the pairs. We then compute the false discovery rate as a function of rank. For `glinternet` and `hierNet`, we obtain a path of solutions and compute the false discovery rate as a function of the number of interactions discovered. The default setting for `hierNet` is to impose weak hierarchy, and we use this except in the cases where the ground truth has strong hierarchy. In these cases, we set `hierNet` to impose strong hierarchy. We also set “diagonal=FALSE” to disable quadratic terms.

We plot the average false discovery rate with standard error bars as a function of the number of predicted interactions in Figure 2. The results are from 100 simulation runs. We see that `glinternet` is competitive with `hierNet` when the truth obeys strong or weak hierarchy, and does better when the truth is anti-hierarchical. This is expected because `hierNet` requires the presence of main effects as a budget for interactions, whereas `glinternet` can still estimate an interaction to be nonzero even though none of its main effects are present. Boosting is not competitive, especially in the anti-hierarchical case. This is because the first split in a tree is effectively looking for a main effect.

Both `glinternet` and `hierNet` perform comparably in these simulations. If all the variables are continuous, there do not seem to be compelling reasons to choose one over the other, apart from the computational advantages of `glinternet` discussed in the next section.

## 6.2 Feasibility

To the best of our knowledge, `hierNet` is the only readily available package for learning interactions among continuous variables in a hierarchical manner. Therefore it is natural to use `hierNet` as a speed benchmark. We generate data in which the ground truth has strong hierarchy as in Section 6.1, but with  $n = 1000$  quantitative responses and  $p = 20, 40, 80, 160, 320, 640$  continuous predictor variables. We set each method to find 10 interactions. While `hierNet` does not allow the user to specify the number of interactions to discover, we get around this by fitting a path of values, then selecting the regularization parameter that corresponds to 10 nonzero estimated interactions. We then refit `hierNet` along a path that terminates with this choice of parameter, and time this run. Both software packages are compiled with the same options. Figure 3 shows the best time recorded for each method over 10 runs. These simulations were timed on a Intel Core-i7 3930K processor.

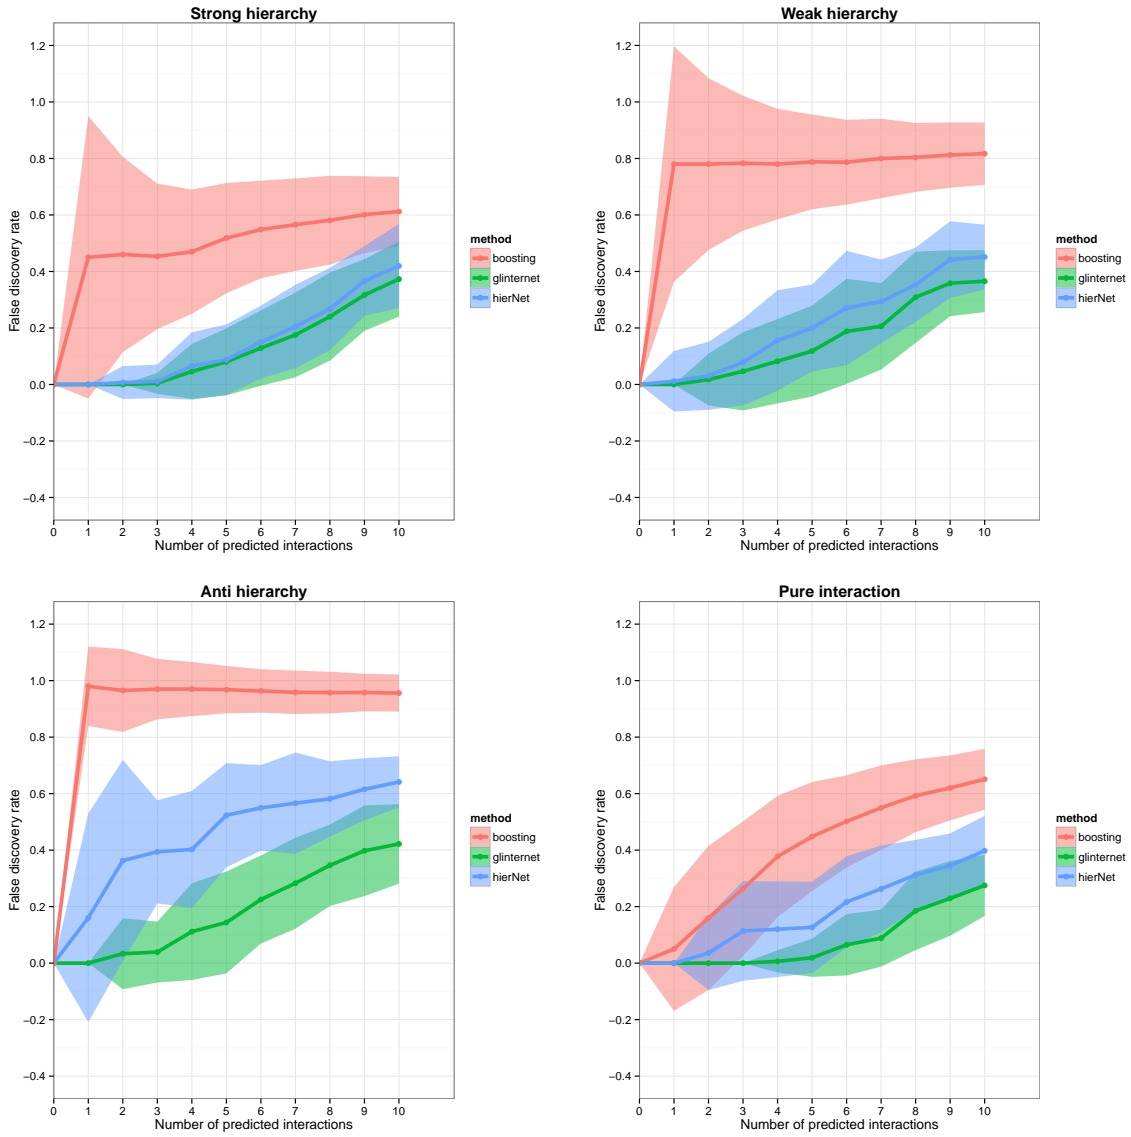


Figure 2: Simulation results for continuous variables: Average false discovery rate and standard errors from 100 simulation runs.






Figure 3: **Left:** Best wallclock time over 10 runs for discovering 10 interactions. **Right:** log scale.

## 7 Real data examples

We compare the performance of `glinternet` on several prediction problems. The competitor methods used are gradient boosting, lasso, ridge regression, and `hierNet` where feasible. In all situations, we determine the number of trees in boosting by first building a model with a large number of trees, typically 5000 or 10000, and then selecting the number that gives the lowest cross-validated error. We use a learning rate of 0.001, and we do not subsample the data since the sample sizes are small in all the cases.

The methods are evaluated on three measures:

1. missclassification error, or 0-1 loss
2. area under the receiver operating characteristic (ROC) curve, or auc
3. cross entropy, given by  $-\frac{1}{n} \sum_{i=1}^n [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$ .

### 7.1 South African heart disease data

The data consists of 462 males from a high risk region for heart disease in South Africa. The task is to predict which subjects had coronary heart disease using risk factors such as cumulative tobacco, blood pressure, and family history of heart disease. We randomly split the data into 362-100 train-test examples, and tuned each method on the training data using 10-fold cross validation before comparing the prediction performance on the held out test data. This splitting process was carried out 20 times, and Figure 4 summarizes the results. The methods are all comparable, with no distinct winner.

### 7.2 Spambase

This is the Spambase data taken from the UCI Machine Learning Repository. There are 4601 binary observations indicating whether an email is spam or non-spam, and 57 integer-valued variables. All the features are log-transformed by  $\log(1 + x)$  before applying the methods. We split the data into a training set consisting of 3065 observations and a test set consisting of 1536 observations. The methods are tuned on the training set using 10-fold cross validation before predicting on the test set. The results are shown in Figure 5.




Figure 4: Performance of methods on 20 train-test splits of the South African heart disease data.




Figure 5: Performance on the Spambase data.

### 7.3 Dorothea

Dorothea is one of the 5 datasets from the NIPS 2003 Feature Learning Challenge, where the goal is to predict if a chemical molecule will bind to a receptor target. There are 100000 binary features that describe three-dimensional properties of the molecules, half of which are probes that have nothing to do with the response. The training, validation, and test sets consist of 800, 350, and 800 observations respectively. More details about how the data were prepared can be found at <http://archive.ics.uci.edu/ml/datasets/Dorothea>.

We run `glinternet` with screening on 1000 main effects, which results in about 100 million candidate interaction pairs. The validation set was used to tune all the methods. We then predict on the test set with the chosen models and submitted the results online for scoring. The best model chosen by `glinternet` made use of 93 features, compared with the 9 features chosen by  $\ell_1$ -penalized logistic regression (lasso). Figure 6 summarizes the performance for each method. We see that `glinternet` has a slight advantage over the lasso, indicating that interactions might be important for this problem. Boosting did not perform well in our false discovery rate simulations, which could be one of the reasons why it does not do well here despite taking interactions into account.

### 7.4 Genome-wide association study

We use the simulated rheumatoid arthritis data (replicate 1) from Problem 3 in Genetic Analysis Workshop 15. Affliction status was determined by a genetic/environmental model to mimic the familial pattern of arthritis; full details can be found in [Miller et al., 2007]. The authors simulated a large population of nuclear families consisting of two parents and two offspring. We are then provided with 1500 randomly chosen families with an affected sibling pair (ASP), and 2,000 unaffected families as a control group. For the control families, we only have data from one randomly chosen sibling. Therefore we also sample one sibling from each of the 1,500 ASPs to obtain 1,500 cases.

There are 9,187 single nucleotide polymorphism (SNP) markers on chromosomes 1 through 22 that are designed to mimic a 10K SNP chip set, and a dense set of 17,820 SNPs on chromosome 6 that approximate the density of a 300K SNP set. Since 210 of the SNPs on chromosome 6 are found in both the dense and non-dense sets, we made sure to include them only once in our analysis. This gives us a total of  $9187 - 210 + 17820 = 26797$  SNPs,




Figure 6: Performance on dorothea

all of which are 3-level categorical variables.

We are also provided with phenotype data, and we include sex, age, smoking history, and the DR alleles from father and mother in our analysis. Sex and smoking history are 2-level categorical variables, while age is continuous. Each DR allele is a 3-level categorical variable, and we combine the father and mother alleles in an unordered way to obtain a 6-level DR variable. In total, we have 26,801 variables and 3,500 training examples.

We run `glinetnet` (without screening) on a grid of values for  $\lambda$  that starts with the empty model. The first two variables found are main effects:

- SNP6\_305
- denseSNP6\_6873

Following that, an interaction `denseSNP6_6881:denseSNP6_6882` gets picked up. We now proceed to analyze this result.

Two of the interactions listed in the answer sheet provided with the data are: locus A with DR, and locus C with DR. There is also a main effect from DR. The closest SNP to the given position of locus A is SNP16\_31 on chromosome 16, so we take this SNP to represent locus A. The given position for locus C corresponds to denseSNP6\_3437, and we use this SNP

for locus C. While it looks like none of the true variables are to be found in the list above, these discovered variables have very strong association with the true variables.

If we fit a linear logistic regression model with our first discovered pair `denseSNP6_6881:denseSNP6_6882`, the main effect `denseSNP6_6882` and the interaction terms are both significant:

	Df	Dev	Resid. Dev	$\mathbb{P}( > \chi^2 )$
NULL			4780.4	
<code>denseSNP6_6881</code>	1	1.61	4778.7	0.20386
<code>denseSNP6_6882</code>	2	1255.68	3523.1	<2e-16
<code>denseSNP6_6881:denseSNP6_6882</code>	1	5.02	3518.0	0.02508

Table 1: Anova for linear model fitted to first interaction term that was discovered.

A  $\chi^2$  test for independence between `denseSNP6_6882` and DR gives a p-value of less than  $1e-15$ , so that `glinternet` has effectively selected an interaction with DR. However, `denseSNP6_6881` has little association with loci A and C. The question then arises as to why we did not find the true interactions with DR. To investigate, we fit a linear logistic regression model separately to each of the two true interaction pairs. In both cases, the main effect DR is significant (p-value <  $1e - 15$ ), but the interaction term is not:

	Df	Dev	Resid. Dev	$\mathbb{P}( > \chi^2 )$
NULL			4780.4	
<code>SNP16_31</code>	2	3.08	4777.3	0.2147
<code>DR</code>	5	2383.39	2393.9	<2e-16
<code>SNP16_31:DR</code>	10	9.56	2384.3	0.4797
NULL			4780.4	
<code>denseSNP6_3437</code>	2	1.30	4779.1	0.5223
<code>DR</code>	5	2384.18	2394.9	<2e-16
<code>denseSNP6_3437:DR</code>	8	5.88	2389.0	0.6604

Table 2: Anova for linear logistic regression done separately on each of the two true interaction terms.

Therefore it is somewhat unsurprising that `glinternet` did not pick these interactions.

This example also illustrates how the the group-lasso penalty in **glinternet** helps in discovering interactions (see Section 3.3). We mentioned above that denseSNP6\_6881:denseSNP6\_6882 is significant if fit by itself in a linear logistic model (Table 1). But if we now fit this interaction *in the presence* of the two main effects SNP6\_305 and denseSNP6\_6873, it is *not* significant:

	Df	Dev	Resid. Dev	$\mathbb{P}( > \chi^2 )$
NULL			4780.4	
SNP6_305	2	2140.18	2640.2	<2e-16
denseSNP6_6873	2	382.61	2257.6	<2e-16
denseSNP6_6881:denseSNP6_6882	4	3.06	2254.5	0.5473

This suggests that fitting the two main effects fully has explained away most of the effect from the interaction. But because **glinternet** regularizes the coefficients of these main effects, they are not fully fit, and this allows **glinternet** to discover the interaction.

The anova analyses above suggest that the true interactions are difficult to find in this GWAS dataset. Despite having to search through a space of about 360 million interaction pairs, **glinternet** was able to find variables that are strongly associated with the truth. This illustrates the difficulty of the interaction-learning problem: even if the computational challenges are met, the statistical issues are perhaps the dominant factor.

## 8 Algorithm details

We describe the algorithm used in **glinternet** for solving the group-lasso optimization problem. Since the algorithm applies to the group-lasso in general and not specifically for learning interactions, we will use  $\mathbf{Y}$  as before to denote the  $n$ -vector of observed responses, but  $\mathbf{X} = [\mathbf{X}_1 \ \mathbf{X}_2 \ \dots \ \mathbf{X}_p]$  will now denote a generic feature matrix whose columns fall into  $p$  groups.

### 8.1 Defining the group penalties $\gamma$

Recall that the group-lasso solves the optimization problem

$$\operatorname{argmin}_{\beta} \mathcal{L}(\mathbf{Y}, \mathbf{X}; \beta) + \lambda \sum_{i=1}^p \gamma_i \|\beta_i\|_2, \quad (74)$$



where  $\mathcal{L}(\mathbf{Y}, \mathbf{X}; \beta)$  is the negative log-likelihood function. This is given by

$$\mathcal{L}(\mathbf{Y}, \mathbf{X}; \beta) = \frac{1}{2n} \|\mathbf{Y} - \mathbf{X}\beta\|_2^2 \quad (75)$$

for squared error loss, and

$$\mathcal{L}(\mathbf{Y}, \mathbf{X}, \beta) = -\frac{1}{n} [\mathbf{Y}^T(\mathbf{X}\beta) - \mathbf{1}^T \log(\mathbf{1} + \exp(\mathbf{X}\beta))] \quad (76)$$

for logistic loss (log and exp are taken component-wise). Each  $\beta_i$  is a vector of coefficients for group  $i$ . When each group consists of only one variable, this reduces to the lasso.

The  $\gamma_i$  allow us to penalize some groups more (or less) than others. We want to choose the  $\gamma_i$  so that if the signal were pure noise, then all the groups are equally likely to be nonzero. Because the quantity  $\|\mathbf{X}_i^T(\mathbf{Y} - \hat{\mathbf{Y}})\|_2$  determines whether the group  $\mathbf{X}_i$  is zero or not (see the KKT conditions (11)), we define  $\gamma_i$  via a null model as follows. Let  $\epsilon \sim (0, I)$ . Then we have

$$\gamma_i^2 = \mathbb{E}\|\mathbf{X}_i^T \epsilon\|_2^2 \quad (77)$$

$$= \text{tr} \mathbf{X}_i^T \mathbf{X}_i \quad (78)$$

$$= \|\mathbf{X}_i\|_F^2. \quad (79)$$

Therefore we take  $\gamma_i = \|\mathbf{X}_i\|_F$ , the frobenius norm of the matrix  $\mathbf{X}_i$ . In the case where the  $\mathbf{X}_i$  are orthonormal matrices with  $p_i$  columns, we recover  $\gamma_i = \sqrt{p_i}$ , which is the value proposed in [Yuan and Lin, 2006]. In our case, the indicator matrices for categorical variables all have frobenius norm equal to  $\sqrt{n}$ , so we can simply take  $\gamma_i = 1$  for all  $i$ . The case where continuous variables are present is not as straightforward, but we can normalize all the groups to have frobenius norm one, which then allows us to take  $\gamma_i = 1$  for  $i = 1, \dots, p$ .

## 8.2 Fitting the group-lasso

Fast iterative soft thresholding (FISTA) [Beck and Teboulle, 2009] is a popular approach for computing the lasso estimates. This is essentially a first order method with Nesterov style acceleration through the use of a momentum factor. Because the group-lasso can be viewed as a more general version of the lasso, it is unsurprising that FISTA can be adapted for the group-lasso with minimal changes. This gives us important advantages:

1. FISTA is a generalized gradient method, so that there is no Hessian involved

2. virtually no change to the algorithm when going from squared error loss to logistic loss
3. gradient computation and parameter updates can be parallelized
4. can take advantage of adaptive momentum restart heuristics.

Adaptive momentum restart was introduced in [O’Donoghue and Candes, 2012] as a scheme to counter the “rippling” behaviour often observed with accelerated gradient methods. They demonstrated that adaptively restarting the momentum factor based on a gradient condition can dramatically speed up the convergence rate of FISTA. The intuition is that we should reset the momentum to zero whenever the gradient at the current step and the momentum point in different directions. Because the restart condition only requires a vector multiplication with the gradient (which has already been computed), the added computational cost is negligible. The FISTA algorithm with adaptive restart is given below.

**Algorithm 1:** FISTA with adaptive restart

**input** : Initialized parameters  $\beta^{(0)}$ , feature matrix  $\mathbf{X}$ , observations  $\mathbf{Y}$ , regularization parameter  $\lambda$ , step size  $s$ .  
**output:**  $\hat{\beta}$

Initialize  $x^{(0)} = \beta^{(0)}$  and  $\rho_0 = 1$ .

**for**  $k = 0, 1, \dots$ , **do**

$g^{(k)} = -\mathbf{X}^T(\mathbf{Y} - \mathbf{X}\beta^{(k)});$
$x^{(k+1)} = \left( \mathbf{1} - \frac{s\lambda}{\ \beta^{(k)} - sg^{(k)}\ _2} \right)_+ (\beta^{(k)} - sg^{(k)});$
$\rho_k = (\beta^{(k)} - x^{(k+1)})^T (x^{(k+1)} - x^{(k)}) > 0 ? 1 : \rho_k;$
$\rho_{k+1} = (1 + \sqrt{1 + 4\rho_k^2})/2;$
$\beta^{(k+1)} = x^{(k+1)} + \frac{\rho_k - 1}{\rho_{k+1}}(x^{(k+1)} - x^{(k)});$

**end**

At each iteration, we take a step in the direction of the gradient with step size  $s$ . We can get an idea of what  $s$  should be by looking at the majorized objective function about a fixed point  $\beta_0$ :

$$M(\beta) = \mathcal{L}(\mathbf{Y}, \mathbf{X}; \beta_0) + (\beta - \beta_0)^T g(\beta_0) + \frac{1}{2s} \|\beta - \beta_0\|_2^2 + \lambda \sum_{i=1}^p \|\beta_i\|_2. \quad (80)$$

Here,  $g(\beta_0)$  is the gradient of the negative log-likelihood  $\mathcal{L}(\mathbf{Y}, \mathbf{X}; \beta)$  evaluated at  $\beta_0$ . Majorization-minimization schemes for convex optimization choose  $s$  sufficiently small so that the LHS of (80) is upper bounded by the RHS. One strategy is to start with a large step size, then backtrack until this condition is satisfied. We use an approach that is mentioned in [Becker et al., 2011] that adaptively initializes the step size with

$$s = \frac{\|\beta^{(k)} - \beta^{(k-1)}\|_2}{\|g_k - g_{k-1}\|_2}. \quad (81)$$

We then backtrack from this initialized value if necessary by multiplying  $s$  with some  $0 < \alpha < 1$ . The interested reader may refer to [Simon et al., 2013] for more details about majorization-minimization schemes for the group-lasso.

## 9 Discussion

We introduced **glinternet**, a method for learning linear interaction models that satisfy strong hierarchy. We demonstrated that the method is comparable with past approaches, but has the added advantage of being able to accommodate both categorical and continuous variables on a larger scale. We illustrated the method with several examples using real and simulated data, and also showed that **glinternet** can be applied to genome wide association studies.

**glinternet** is available on CRAN as a package for the statistical software R. The user has the option of using multiple cores/cpus, but this requires that the package be compiled with OpenMP enabled. The user may also benefit from compiling with a higher level of optimization (such as O3 for gcc). For example, all our results were obtained by compiling with the following command (same command used for **hierNet**)

```
icc -std=gnu99 -xhost -fp-model fast=2 -no-prec-div -no-prec-sqrt -ip -O3 -restrict
```

## References

- [Beck and Teboulle, 2009] Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Img. Sci.*, 2(1):183–202.
- [Becker et al., 2011] Becker, S., Candes, E., and Grant, M. (2011). Templates for convex cone problems with applications to sparse signal recovery. *Mathematical Programming Computation*, 3(3):165–218.
- [Bien et al, 2013] Bien et al, J. (2013). A lasso for hierarchical interactions. *The Annals of Statistics*.
- [Chen et al., 2011] Chen, C., Schwender, H., Keith, J., Nunkesser, R., Mengersen, K., and Macrossan, P. (2011). Methods for identifying snp interactions: A review on variations of logic regression, random forest and bayesian logistic regression. *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, 8(6):1580 –1591.
- [Freund and Schapire, 1995] Freund, Y. and Schapire, R. E. (1995). A decision-theoretic generalization of on-line learning and an application to boosting. In *Proceedings of the Second European Conference on Computational Learning Theory, EuroCOLT '95*, pages 23–37, London, UK, UK. Springer-Verlag.
- [Friedman, 2000] Friedman, J. H. (2000). Greedy Function Approximation: A Gradient Boosting Machine. In *Annals of Statistics*, volume 29, pages 1189–1232.
- [Jacob et al., 2009] Jacob, L., Obozinski, G., and Vert, J.-P. (2009). Group lasso with overlap and graph lasso. In *Proceeding of the 26th International Conference on Machine Learning, Montreal, Canada*.
- [Koren, 2009] Koren, Y. (2009). Collaborative filtering with temporal dynamics. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD'09)*.
- [Miller et al., 2007] Miller, M., Li, M., Lind, G., and Jang, S.-Y. (2007). Problem 3: Simulated rheumatoid arthritis data. *BMC Proceedings 1 (Suppl1):S4*.
- [O’Donoghue and Candes, 2012] O’Donoghue, B. and Candes, E. (2012). Adaptive restart for accelerated gradient schemes. *Foundations of Computational Mathematics*.

- [Ruczinski I, 2003] Ruczinski I, Kooperberg C, L. M. (2003). Logic regression. *Journal of Computational and Graphical Statistics*, 12(3):475–511.
- [Scheffe, 1959] Scheffe, L. (1959). *The Analysis of Variance*. Wiley.
- [Schwender and Ickstadt, 2008] Schwender, H. and Ickstadt, K. (2008). Identification of snp interactions using logic regression. *Biostatistics*, 9(1):187–198.
- [Simon et al., 2013] Simon, N., Friedman, J., Hastie, T., and Tibshirani, R. (2013). A sparse-group lasso. *Journal of Computational and Graphical Statistics*, 22(2):231–245.
- [Tibshirani, 1996] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society - Series B: Statistical Methodology*, 58(1):267–288.
- [Tibshirani et al, 2012] Tibshirani et al, R. (2012). Strong rules for discarding predictors in lasso-type problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- [Yuan and Lin, 2006] Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society - Series B: Statistical Methodology*, 68(1):49–67.
- [Zhao et al., 2009] Zhao, P., Rocha, G., and Yu, B. (2009). Grouped and hierarchical model selection through composite absolute penalties. *The Annals of Statistics*, 37 (6A):3468–3497.