Multivariate GWAS: Generalized Linear Models, Prior Weights, and Double Sparsity

Benjamin B. Chu[1], Kevin L. Keys[2], Janet S. Sinsheimer[1,3,*], Kenneth Lange[1,3†]

[1]Department of Computational Medicine, David Geffen School of Medicine at UCLA, Los Angeles, USA
[2]Department of Medicine, University of California, San Francisco, USA
[3]Department of Human Genetics, David Geffen School of Medicine at UCLA, Los Angeles, USA

keywords: GWAS; multivariate regression; iterative hard thresholding; generalized linear models

1 Abstract

Background: Consecutive testing of single nucleotide polymorphisms (SNPs) is usually employed to identify genetic variants associated with complex traits. Ideally one should model all covariates in unison, but most existing analysis methods for genome-wide association studies (GWAS) perform only univariate regression.

Results: We extend and efficiently implement iterative hard thresholding (IHT) for multivariate regression. Our extensions accommodate generalized linear models (GLMs), prior information on genetic variants, and grouping of variants. In our simulations, IHT recovers up to 30% more true predictors than SNP-by-SNP association testing, and exhibits a 2 to 3 orders of magnitude decrease in false positive rates compared to lasso regression. These advantages capitalize on IHT's ability to recover unbiased coefficient estimates. We also apply IHT to the Northern Finland Birth Cohort of 1966 and find that IHT recovers plausible variants associated with HDL and LDL.

Conclusions: Our real data analysis and simulation studies suggest that IHT can (a) recover highly correlated predictors, (b) avoid over-fitting, (c) deliver better true positive and false positive rates than either marginal testing or lasso regression, (d) recover unbiased regression coefficients, and (e) exploit prior information and group-sparsity. Although these advances are studied for GWAS inference, our extensions are pertinent to other regression problems with large numbers of predictors.

*Corresponding author. Email: jsinshei@ucla.edu
†Corresponding author. Email: klange@ucla.edu

2 Introduction

In genome-wide association studies (GWAS), modern genotyping technology coupled with imputation algorithms can produce an $n \times p$ genotype matrix $\mathbf{X}$ with $n \approx 10^6$ subjects and $p \approx 10^7$ genetic predictors (8; 35). Data sets of this size require hundreds of gigabytes of disk space to store in compressed form. Decompressing data to floating point numbers for statistical analyses leads to matrices too large to fit into standard computer memory. The computational burden of dealing with massive GWAS datasets limits statistical analysis and interpretation. This paper discusses and extends a class of algorithms capable of meeting the challenge of multivariate regression with modern GWAS data scales.

Traditionally, GWAS analysis has focused on SNP-by-SNP (single nucleotide polymorphism) association testing (8; 7), with a p-value computed for each SNP via linear regression. This approach enjoys the advantages of simplicity, interpretability, and a low computational complexity of $\mathcal{O}(np)$. Furthermore, since the genotype matrix can be streamed column by column, marginal linear regressions make efficient use of memory. Some authors further increase association power by reframing GWAS as a linear mixed model problem and proceeding with variance component selection (17; 23). These advances remain within the scope of marginal analysis.

Despite their numerous successes (35), marginal regression is less than ideal for GWAS. Multivariate statistical methods can in principle treat all SNPs simultaneously. This approach captures the biology behind GWAS more realistically because traits are usually determined by multiple SNPs acting in unison. Marginal regression selects associated SNPs one by one based on a pre-set threshold. Given the stringency of the p-value threshold, marginal regression can miss many causal SNPs with low effect sizes. As a result, heritability is underestimated. When $p \gg n$, one usually assumes that the number of variants k associated with a complex trait is much less than n. If this is true, we can expect multivariate regression to perform better because it a) offers better outlier detection and better prediction, b) accounts for the correlations among SNPs, and c) allows investigators to model interactions. Of course, these advantages are predicated on finding the truly associated SNPs.

Adding penalties to the loss function is one way of achieving parsimony in regression. The lasso (33; 34) is the most popular model selection device in current use. The lasso model selects non-zero parameters by minimizing the criterion

$$f(\boldsymbol{\beta}) \quad = \quad \ell(\boldsymbol{\beta}) + \lambda \|\boldsymbol{\beta}\|_1,$$

where $\ell(\boldsymbol{\beta})$ is a convex loss, λ is a sparsity tuning constant, and $\|\boldsymbol{\beta}\|_1 = \sum_j |\beta_j|$ is the ℓ_1 norm of the parameters. The lasso has the virtues of preserving convexity and driving most parameter estimates to 0. Minimization can be conducted efficiently via cyclic coordinate descent (13; 37). The magnitude of the nonzero tuning constant λ determines the number of predictors selected.

Despite its widespread use, the lasso penalty has some drawbacks. First, the ℓ_1 penalty tends to shrink parameters toward 0, sometimes severely so. Second, λ must be tuned to achieve a given model size. Third, λ is chosen by cross-validation, a costly procedure. Fourth and most importantly, the shrinkage caused by the penalty tends to encourage too many false positives to enter the model ultimately identified by cross-validation.

Inflated false positive rates can be mitigated by substituting nonconvex penalties for the ℓ_1 penalty. For

example, the minimax concave penalty (MCP) (42)

$$\lambda p(\beta_j) \quad = \quad \lambda \int_0^{|\beta_j|} \left(1 - \frac{s}{\lambda\gamma}\right)_+ ds$$

starts out at $\beta_j = 0$ with slope λ and gradually transitions to a slope of 0 at $\beta_j = \lambda\gamma$. With minor adjustments, the coordinate descent algorithm for the lasso carries over to MCP penalized regression (6; 26). Model selection is achieved without severe shrinkage, and inference in GWAS improves (18). However, in our experience the false negative rate increases under cross validation (19). A second remedy for the lasso, stability selection, weeds out false positives by looking for consistent predictor selection across random halves of the data (1; 29).

In contrast, iterative hard thresholding (IHT) minimizes a loss $\ell(\boldsymbol{\beta})$ subject to the nonconvex sparsity constraint $\|\boldsymbol{\beta}\|_0 \leq k$, where $\|\boldsymbol{\beta}\|_0$ counts the number of non-zero components of $\boldsymbol{\beta}$ (2; 3; 5). Figure 1 explains graphically how the ℓ_0 penalty reduces the bias of the selected parameters. The reduced bias in IHT regression leads to more accurately estimated effect sizes. For GWAS, the sparsity model-size constant k also has a simpler and more intuitive interpretation than the lasso tuning constant λ. Finally, both false positive and false negative rates are well controlled. Balanced against these advantages is the loss of convexity in optimization and concomitant loss of computational efficiency. In practice, the computational barriers are surmountable and are compensated by the excellent results delivered by IHT in high-dimensional regression problems such as multivariate GWAS.

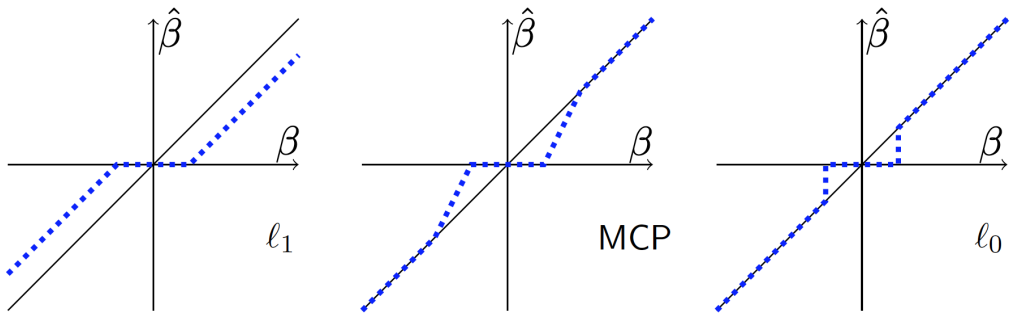

Figure 1: The ℓ_0 quasinorm of IHT enforces sparsity without shrinkage. The estimated effect size ($\hat{\beta}$) is plotted against its true value (β) for ℓ_1, MPC, and ℓ_0 penalties.

This article has four interrelated goals. First, we extend IHT to generalized linear models. These models encompass most of applied statistics. Previous IHT algorithms focused on normal or logistic sparse regression scenarios. Our software performs sparse regression under more exotic Poisson and negative binomial response distributions and can be easily extended to other GLM distributions as needed. The key to our extension is the derivation of a nearly optimal step size s for improving the loglikelihood at each iteration. Second, we introduce doubly-sparse regression to IHT. Previous authors have considered group sparsity (39). The latter tactic limits the number of groups selected. It is also useful to limit the number of predictors selected per group. Double sparsity strikes a compromise that encourages selection of correlated causative variants in linkage disequilibrium (LD). Third, we demonstrate how to incorporate SNP weights in rare variant discovery. Our simple and interpretable weighting scheme directly introduces prior knowledge into sparse projection. This allows one to favor predictors whose association to the response is supported by external evidence. Fourth, we provide scalable, open source, and user friendly software for IHT. On a modern laptop, our code can handle datasets with 10^5 subjects and half a million SNPs. This is provided in a Julia (4) package

called `MendelIHT.jl`, interfacing with the OpenMendel umbrella (45) and JuliaStats's Distribution and GLM packages (22).

3 Model Development

This section sketches our extensions of iterative hard thresholding (IHT).

3.1 IHT Background

IHT was originally formulated for sparse signal reconstruction, which is framed as sparse linear least squares regression. In classical linear regression, we are given an $n \times p$ design matrix $\mathbf{X}$ and a corresponding n-component response vector $\mathbf{y}$. We then postulate that $\mathbf{y}$ has mean $\mathrm{E}(\mathbf{y}) = \mathbf{X}\boldsymbol{\beta}$ and that the residual vector $\mathbf{y} - \mathbf{X}\boldsymbol{\beta}$ has independent Gaussian components with a common variance. The parameter (regression coefficient) vector $\boldsymbol{\beta}$ is estimated by minimizing the sum of squares $f(\boldsymbol{\beta}) = \frac{1}{2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2$. The solution to this problem is known as the ordinary least squares estimator and can be written explicitly as $\hat{\boldsymbol{\beta}} = (\mathbf{X}^t\mathbf{X})^{-1}\mathbf{X}^t\mathbf{y}$, provided the problem is overdetermined (n > p). This paradigm breaks down in the big data regime $n \ll p$, where the parameter vector $\boldsymbol{\beta}$ is underdetermined. In the spirit of parsimony, IHT seeks a sparse version of $\boldsymbol{\beta}$ that gives a good fit to the data. This is accomplished by minimizing $f(\boldsymbol{\beta})$ subject to $\|\boldsymbol{\beta}\|_0 \leq k$ for a small value of k, where $\|\cdot\|_0$ counts the number of nonzero entries of a vector. IHT abandons the explicit formula for $\hat{\boldsymbol{\beta}}$ because it fails to respect sparsity and involves the numerically intractable matrix inverse $(\mathbf{X}^t\mathbf{X})^{-1}$.

IHT combines three core ideas. The first is steepest descent. Elementary calculus tells us that the negative gradient $-\nabla f(\mathbf{x})$ is the direction of steepest descent of $f(\boldsymbol{\beta})$ at $\mathbf{x}$. First-order optimization methods like IHT define the next iterate in minimization by the formula $\boldsymbol{\beta}_{n+1} = \boldsymbol{\beta}_n + s_n\mathbf{v}_n$, where $\mathbf{v}_n = -\nabla f(\boldsymbol{\beta}_n)$ and $s_n > 0$ is some optimally chosen step size. In the case of linear regression $-\nabla f(\boldsymbol{\beta}) = \mathbf{X}^t(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})$. To reduce the error at each iteration, the optimal step size s_n can be selected by minimizing the second-order Taylor expansion

$$\begin{aligned} f(\boldsymbol{\beta}_n + s_n\mathbf{v}_n) &= f(\boldsymbol{\beta}_n) + s_n\nabla f(\boldsymbol{\beta}_n)^t\mathbf{v}_n + \frac{s_n^2}{2}\mathbf{v}_n^t d^2 f(\boldsymbol{\beta}_n)\mathbf{v}_n \\ &= f(\boldsymbol{\beta}_n) - s_n\|\nabla f(\boldsymbol{\beta}_n)\|_2^2 + \frac{s_n^2}{2}\nabla f(\boldsymbol{\beta}_n)^t d^2 f(\boldsymbol{\beta}_n)\nabla f(\boldsymbol{\beta}_n) \end{aligned}$$

with respect to s_n. Here $d^2 f(\boldsymbol{\beta}) = \mathbf{X}^t\mathbf{X}$ is the Hessian matrix of second partial derivatives. Because $f(\boldsymbol{\beta})$ is quadratic, the expansion is exact. Its minimum occurs at the step size

$$s_n = \frac{\|\nabla f(\boldsymbol{\beta}_n)\|_2^2}{\nabla f(\boldsymbol{\beta}_n)^t d^2 f(\boldsymbol{\beta}_n)\nabla f(\boldsymbol{\beta}_n)}. \tag{3.1}$$

This formula summarizes the second core idea.

The third component of IHT involves projecting the steepest descent update $\boldsymbol{\beta}_n + s_n\mathbf{v}_n$ onto the sparsity set $S_k = \{\boldsymbol{\beta} : \|\boldsymbol{\beta}\|_0 \leq k\}$. The relevant projection operator $P_{S_k}(\boldsymbol{\beta})$ sets all but the k largest entries of $\boldsymbol{\beta}$ in magnitude to 0. In summary, IHT updates the parameter vector $\boldsymbol{\beta}$ according to the recipe

$$\boldsymbol{\beta}_{n+1} = P_{S_k}(\boldsymbol{\beta}_n - s_n\nabla f(\boldsymbol{\beta}_n))$$

Family	Mean Domain	Var(y)	$g(s)$
Normal	$\mathbb{R}$	ϕ^2	1
Poisson	$[0,\infty)$	μ	e^s
Bernoulli	$[0,1]$	$\mu(1-\mu)$	$\frac{e^s}{1+e^s}$
Gamma	$[0,\infty)$	$\mu^2\phi$	s^{-1}
Inverse Gaussian	$[0,\infty)$	$\mu^3\phi$	$s^{-1/2}$
Negative Binomial	$[0,\infty)$	$\mu(\mu\phi+1)$	e^s

Table 1: Summary of mean domains and variances for common exponential distributions. In GLM, $\boldsymbol{\mu} = g(\mathbf{x}^t\boldsymbol{\beta})$ denotes the mean, $s = \mathbf{x}^t\boldsymbol{\beta}$ the linear responses, g is the inverse link function, and ϕ the dispersion. Except for the negative binomial, all inverse links are canonical.

with the step size given by formula (3.1).

An optional debiasing step can be added to improve parameter estimates. This involves replacing $\boldsymbol{\beta}_{n+1}$ by the exact minimum point of $f(\boldsymbol{\beta})$ in the subspace defined by the support $\{j : \boldsymbol{\beta}_{n+1,j} \neq 0\}$ of $\boldsymbol{\beta}_{n+1}$. Debiasing is efficient because it solves a low-dimensional problem. Several versions of hard-thresholding algorithms have been proposed in the signal processing literature. The first of these, NIHT (5), omits debaising. The rest, HTP(12), GraHTP (40), and CoSaMp (31) offer debiasing.

3.2 IHT for Generalized Linear Models

A generalized linear model (GLM) involves responses y following a natural exponential distribution with density in the canonical form

$$f(y \mid \theta, \phi) \quad = \quad \exp\left[\frac{y\theta - b(\theta)}{a(\phi)} + c(y,\phi)\right],$$

where y is the data, θ is the natural parameter, $\phi > 0$ is the scale (dispersion), and $a(\phi)$, $b(\theta)$, and $c(y,\phi)$ are known functions which vary depending on the distribution (11; 27). Simple calculations show that y has mean $\mu = b'(\theta)$ and variance $\sigma^2 = b''(\theta)a(\phi)$; accordingly, σ^2 is a function of μ. Table 1 summarizes the mean domains and variances of a few common exponential families. Covariates enter GLM modeling through an inverse link representation $\mu = g(\mathbf{x}^t\boldsymbol{\beta})$, where $\mathbf{x}$ is a vector of covariates (predictors) and $\boldsymbol{\beta}$ is vector of regression coefficients (parameters). In statistical practice, data arrive as a sample of independent responses $y_1, \ldots, y_m$ with different covariate vectors $\mathbf{x}_1, \ldots, \mathbf{x}_m$. To put each predictor on an equal footing, each should be standardized to have mean 0 and variance 1. Including an additional intercept term is standard practice.

If we assemble a design matrix $\mathbf{X}$ by stacking the row vectors $\mathbf{x}_i^t$, then we can calculate the loglikelihood,

score, and expected information (11; 20; 27; 38)

$$
\begin{aligned}
L(\boldsymbol{\beta}) &= \sum_{i=1}^{n}\left[\frac{y_i\theta_i - b_i(\theta_i)}{a_i(\phi_i)} + c(y_i,\phi_i)\right] \\
\nabla L(\boldsymbol{\beta}) &= \sum_{i=1}^{n}(y_i-\mu_i)\frac{g'(\mathbf{x}_i^t\boldsymbol{\beta})}{\sigma_i^2}\mathbf{x}_i \;=\; \mathbf{X}^t\mathbf{W}_1(\mathbf{y}-\boldsymbol{\mu}) \\
J(\boldsymbol{\beta}) &= \sum_{i=1}^{n}\frac{1}{\sigma_i^2}g'(\mathbf{x}_i^t\boldsymbol{\beta})^2\mathbf{x}_i\mathbf{x}_i^t \;=\; \mathbf{X}^t\mathbf{W}_2\mathbf{X},
\end{aligned}
\tag{3.2}
$$

where $\mathbf{W}_1$ and $\mathbf{W}_2$ are two diagonal matrices. The second has positive diagonal entries; they coincide under the identity inverse link $g(s) = s$. In the generalized linear model version of IHT, we maximize $L(\boldsymbol{\beta})$ (equivalent to minimizing $f(\boldsymbol{\beta}) = -L(\boldsymbol{\beta})$) and substitute the expected information $J(\boldsymbol{\beta}_n) = \mathrm{E}[-d^2L(\boldsymbol{\beta}_n)]$ for $d^2f(\boldsymbol{\beta}_n)$ in formula (3.1). This translates into the following step size in GLM estimation:

$$
s_n \;=\; \frac{\|\nabla L(\boldsymbol{\beta}_n)\|_2^2}{\nabla L(\boldsymbol{\beta}_n)^t J(\boldsymbol{\beta}_n)\nabla L(\boldsymbol{\beta}_n)}. \tag{3.3}
$$

This substitution is a key ingredient of our extended IHT. It simplifies computations and guarantees that the step size is nonnegative.

3.3 Doubly Sparse Projections

The effectiveness of group sparsity in penalized regression has been demonstrated in general (28; 14) and for GWAS (44) in particular. Group IHT was introduced by Yang et al (39), who enforces a sparse number of groups but do not enforce within-group sparsity. In GWAS, two causative SNPs can be highly correlated with each other due to linkage disequilibrium (LD). Here we discuss how to carry out a doubly-sparse projection that enforces both within- and between-group sparsity. This tactic encourages the detection of such correlated SNPs.

Suppose we divide the SNPs of a study into a collection G of nonoverlapping groups. Given a parameter vector $\boldsymbol{\beta}$ and a group $g \in G$, let $\boldsymbol{\beta}_g$ denote the components of $\boldsymbol{\beta}$ corresponding to the SNPs in g. Now suppose we want to select at most j groups and at most k SNPs per group. In projecting $\boldsymbol{\beta}$, the component β_i is untouched for a selected SNP i. For an unselected SNP, β_i is reset to 0. By analogy with our earlier discussion, we can define a sparsity projection operator $P_{kg}(\boldsymbol{\beta}_g)$ for each group g; $P_{kg}(\boldsymbol{\beta}_g)$ selects the k most prominent SNPs in g. The potential reduction in the squared distance offered by group g is $r_g = \|\boldsymbol{\beta}_g\|_2^2 - \|P_{kg}(\boldsymbol{\beta}_g)\|_2^2$. The j selected groups are determined by selecting the j largest values of r_g. In Algorithm 1, we write $\boldsymbol{\beta} \in S_{jk}$ whenever $\boldsymbol{\beta}$ has at most j active groups with at most k predictors per group.

3.4 Prior weights in IHT

Zhou et al. (44) treat prior weights in penalized GWAS. Before calculating the lasso penalty, they multiply each component of the parameter vector $\boldsymbol{\beta}$ by a positive weight w_i. We can do the same in IHT before projection. Thus, instead of projecting the steepest descent step $\boldsymbol{\beta} = \boldsymbol{\beta}_n + s_n\mathbf{v}_n$, we project the Hadamard (pointwise) product $\mathbf{w} \circ \boldsymbol{\beta}$ of $\boldsymbol{\beta}$ with a weight vector $\mathbf{w}$. This produces a vector with a sensible support S. The next iterate $\boldsymbol{\beta}_{n+1}$ is defined to have support S and to be equal to $\boldsymbol{\beta}_n + s_n\mathbf{v}_n$ on S. The simplest scheme for

choosing nonconstant weights for GWAS relies on minor allele frequencies. Following Zhou et al. (43), we assign SNP i with minor allele frequency p_i the weight $w_i = 1/\sqrt{2p_i(1-p_i)}$. Thus, rare SNPs are weighted more heavily than common SNPs. As we show in results section, users can also assign weights based on pathway and gene information.

de Lamare et al. (10) incorporate prior weights into IHT by adding an element-wise logarithm of a weight vector $\mathbf{q}$ before projection. The weight vector $\mathbf{q}$ is updated iteratively and requires two additional tuning constants that in practice are only obtained through cross validation. Our weighting scheme is simpler, more computationally efficient, and more interpretable.

3.5 Algorithm Summary

The final algorithm combining doubly sparse projections, prior weight scaling, and debiasing is summarized in Algorithm 1.

Algorithm 1: Generalized Iterative hard-thresholding

Input : design matrix $\mathbf{X}$, response vector $\mathbf{y}$, membership vector $\mathbf{g}$, weight vector $\mathbf{w}$, max number of groups j, max predictors per group k.

1 **Initialize:** $\boldsymbol{\beta} \equiv \mathbf{0}$.
2 **while not converged do**
3 **Calculate:** score $= v$, Fisher information matrix $= J$, and step size $= s = \frac{v^T v}{v^T J v}$
4 **Ascent direction with scaling:** $\tilde{\boldsymbol{\beta}} = \mathbf{w} \circ (\boldsymbol{\beta}_n + sv)$
5 **Project to sparsity:** $\tilde{\boldsymbol{\beta}} = P_{S_{jk}}\left(\tilde{\boldsymbol{\beta}}\right) ./\mathbf{w}$ (where $./$ is elementwise division)
6 **while** $L(\tilde{\boldsymbol{\beta}}) \leq L(\boldsymbol{\beta}_n)$**, backtrack** ≤ 5 **do**
7 $s = s/2$
8 Redo lines 4 to 5
9 **end**
10 **(Optional) Debias:** Let $F = \text{supp}(\tilde{\boldsymbol{\beta}})$, compute $\hat{\boldsymbol{\beta}} = \text{argmax}_{\{\boldsymbol{\beta}:\boldsymbol{\beta} \text{ restricted to } F\}} L(\boldsymbol{\beta})$
11 **Accept proposal:** $\boldsymbol{\beta}_{n+1} = \hat{\boldsymbol{\beta}}$
12 **end**

Output: $\boldsymbol{\beta}$ with j active groups and k active predictors per group

4 Results

4.1 Scalability of IHT

To test the scalability of our implementation, we ran IHT on $p = 10^6$ SNPs for sample sizes $n = 10,000, 20,000, \ldots, 120,000$ with five independent replicates per n. All simulations rely on a true sparsity level of $k = 10$. Based on an Intel-E5-2670 machine with 63GB of RAM and a single 3.3GHz processor, Figure 2 plots the IHT median CPU time per iteration, median iterations to convergence, and median memory usage under Gaussian, logistic, Poisson, and negative binomial models. The largest matrix simulated here is 30GB in size and can still fit into our personal computer's memory. Of course, it is possible to test even larger sample sizes using cloud or cluster resources, which are often needed in practice.

The formation of the vector $\boldsymbol{\mu}$ of predicted values requires only a limited number of nonzero regression coefficients. Consequently, the computational complexity of this phase of IHT is relatively light. In contrast, calculation of the Fisher score (gradient) and information (expected negative Hessian) depend on the entire genotype matrix $\mathbf{X}$. Fortunately, each of the np entries of $\mathbf{X}$ can be compressed to 2 bits. Figure 2b and d show that IHT memory demands beyond storing $\mathbf{X}$ never exceeded a few gigabytes. Figure 2a and c show that IHT run time per iteration increases linearly in problem size n. Debiasing increases run time per iteration only slightly. Except for negative binomial responses, debiasing is effective in reducing the number of iterations required for convergence and hence overall run time.

(a) Speed per iteration without debiasing

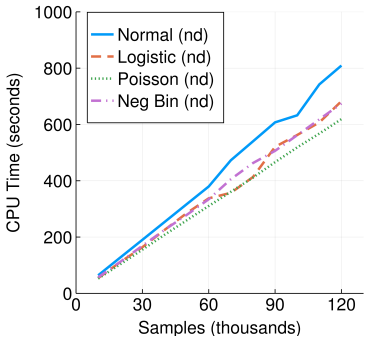

(b) Memory usage without debiasing

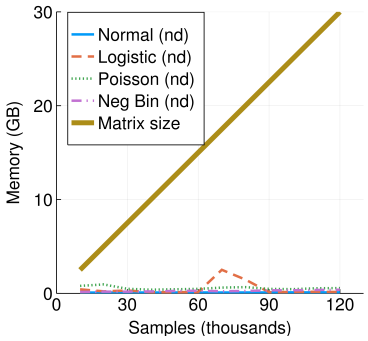

(c) Median iterations until convergence.

	Without debiasing	With Debiasing
Normal	4.0 (1.0)	2.0 (0.0)
Logistic	10.0 (5.5)	2.5 (2.25)
Poisson	47.5 (9.75)	33.0 (5.75)
Neg Bin	9.0 (0.25)	15.0 (1.5)

( ) = interquartile range

(d) Speed per iteration with debiasing

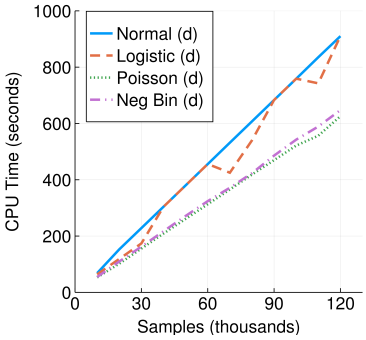

(e) Memory usage with debiasing

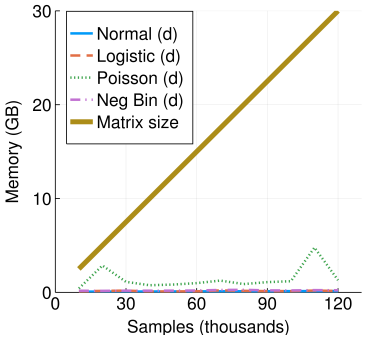

Figure 2: (a, d) Time per iteration scales linearly with data size. Speed is measured for compressed genotype files. On uncompressed data, all responses are roughly 10 times faster. (b, e) Memory usage scales as $\sim 2np$ bits. Note memory for each response are usages in addition to loading the genotype matrix. Uncompressed data requires 32 times more memory. (c) Debiasing reduces median iterations until convergence for all but negative binomial regression. Benchmarks were carried out on 10^6 SNPs and sample sizes ranging from 10,000 to 120,000. Hence, the largest matrix here requires 30GB and can still fit into personal computer memories.

4.2 Cross Validation in Model Selection

In actual studies, the true number of genetic predictors k_{true} is unknown. This section investigates how q-fold cross-validation can determine the best model size on simulated data. Under normal, logistic, Poisson, and negative binomial models, we considered 30 different combinations of $\mathbf{X}$, $\mathbf{y}$, and $\boldsymbol{\beta}_{\text{true}}$ with $k_{\text{true}} = 10$,

$n = 5000$ samples, and $p = 50,000$ SNPs fixed in all replicates. On these data sets we conducted 5-fold cross validation across 20 model sizes k ranging from 1 to 20. Figure 3 plots deviance residuals for each of the four GLM responses (mean squared error in the case of normal responses) and the best estimate $\hat{k}$ of k_{true}.

Figure 3 shows that k_{true} can be effectively recovered by cross validation. In general, prediction error starts off high where the proposed sparsity level k severely underestimates k_{true} and plateaus when k_{true} is reached (Figure 3a-d). Furthermore, the estimated sparsity $\hat{k}$ for each run is narrowly centered around $k_{\text{true}} = 10$ (Figure 3e-f). In fact, $|\hat{k} - k_{\text{true}}| \leq 4$ always holds. When $\hat{k}$ exceeds k_{true}, the estimated regression coefficients for the false predictors tend to be very small. In other words, IHT is robust to overfitting, in contrast to lasso penalized regression.

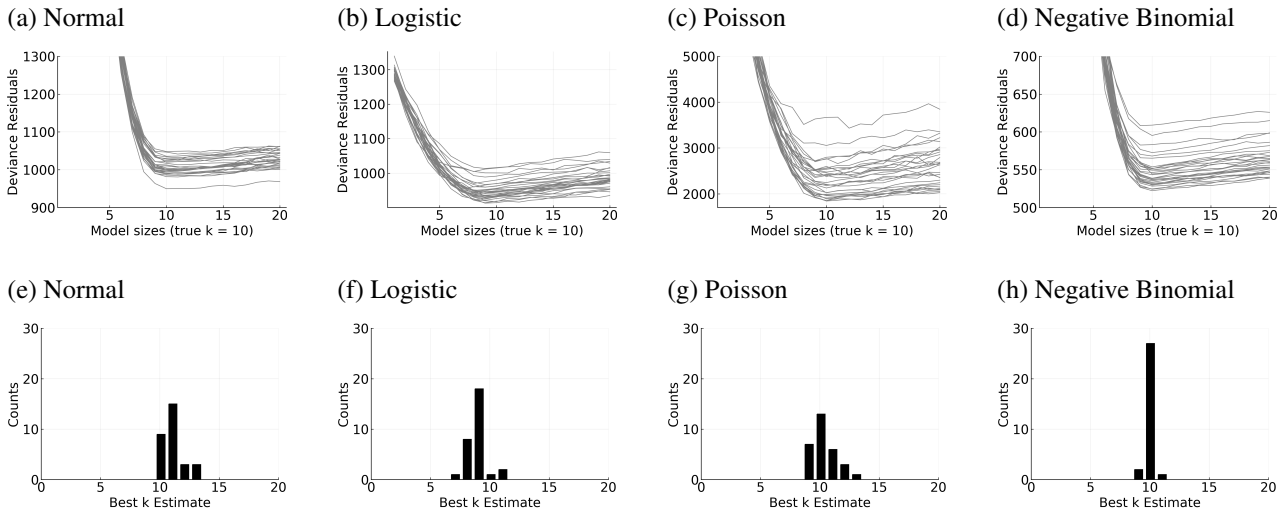

Figure 3: Cross validation is capable of identifying the true model size k_{true}. (a-d) Deviance residuals are minimized when the estimated model size $\hat{k} \approx k_{\text{true}}$. Each line represents 1 simulation. (e-h) $\hat{k}$ is narrowly spread around $k_{\text{true}} = 10$.

4.3 Comparing IHT to Lasso and Marginal Tests in Model Selection

Comparison of the true positive and false positive rates of IHT and its main competitors is revealing. For lasso regression we use the `glmnet` implementation of cyclic coordinate descent (13; 36; 37) (v2.0-16 implemented in `R 3.5.2`); for marginal testing we use the beta version of `MendelGWAS` (45). As explained later, Poisson regression is supplemented by zero-inflated Poisson regression implemented under the `pscl` (41) (v1.5.2) package of `R`. Unfortunately, `glmnet` does not accommodate negative binomial regression. Because both `glmnet` and `pscl` operate on floating point numbers, we limit our comparisons to small problems with 1000 subjects, 10,000 SNPs, 30 replicates, and $k = 10$ causal SNPs. IHT performs model selection by 3-fold cross validation across model sizes ranging from 1 to 50. This range is generous enough to cover the models selected by lasso regression. We adjust for multiple testing in the marginal case test by applying a p-value cutoff of 5×10^{-6}.

	Normal	Logistic	Poisson	Neg Bin
IHT TP	8.84	6.28	7.2	9.0
IHT FP	0.02	0.1	1.28	0.98
Lasso TP	9.52	8.16	9.28	NA
Lasso FP	31.26	45.76	102.24	NA
Marginal TP	7.18	5.76	9.04 (5.94*)	5.98
Marginal FP	0.06	0.02	1527.9 (0.0*)	0.0

Table 2: IHT is superior to lasso and marginal regressions in balancing false positive rates and true positive rates. TP indicates true positives and FP indicates false positives. Displayed rates average 30 independent runs with $k = 10$ causal SNPs. *The parenthesized marginal Poisson result reflects zero-inflated Poisson regression.

Table 2 demonstrates that IHT achieves the best balance between maximizing true positives and minimizing false positives. IHT finds more true positives than marginal testing and almost as many as lasso regression. IHT also finds far fewer false positives than lasso regression. Poisson regression is exceptional in yielding an excessive number of false positives in marginal testing. A similar but less extreme trend is observed for lasso regression. The marginal false positive rate is reduced by switching to zero-inflated Poisson regression, which allows for overdispersion. Interestingly IHT rescues the Poisson model by accurately capturing the simultaneous impact of multiple predictors.

4.4 Reconstruction Quality for GWAS Data

Table 3 demonstrates that IHT estimates show little bias. However, as the magnitude of β_{true} falls, the estimates do show an upward absolute bias consistent with the winner's curse phenomenon. These trends hold with or without implementation of the debiasing procedure described earlier. The proportion of variance explained is approximately the same under both scenarios. The results displayed in Table 3 reflect $n = 10,000$ subjects, $p = 100,000$ SNPs, 100 replicates, and a sparsity level k fixed at it true value $k_{\text{true}} = 10$. To avoid data sets with monomorphic SNPs, the minimum minor allele frequency (maf) is set at 0.05. Readers interested in comparing these parameter estimates with those delivered by lasso regression can visit our Github page. As expected, the lasso estimates exhibit strong shrinkage.

β_{true}	β_{Normal}	β_{Logistic}	β_{Poisson}	$\beta_{\text{Neg Bin}}$
.50	$.500 \pm .010$	$.504 \pm .022$	$.473 \pm .046$	$.476 \pm .045$
.25	$.250 \pm .009$	$.252 \pm .021$	$.236 \pm .026$	$.238 \pm .024$
.10	$.097 \pm .009$	$.108 \pm .015$	$.096 \pm .011$	$.096 \pm .012$
.05	$.053 \pm .008$	$.090 \pm .004$	$.051 \pm .008$	$.054 \pm .006$
.03	$.046 \pm .004$	NA	$.045 \pm .005$	$.049 \pm .007$

Table 3: IHT recovers nearly unbiased coefficient estimates. Displayed coefficients are average fitted valued $\pm$ one standard error for the discovered predictors.

4.5 Correlated Covariates and Doubly Sparse Projections

Next we study how well IHT works on correlated data and whether doubly-sparse projection can enhance model selection. Table 4 shows that, in the presence of extensive LD, IHT performs reasonably well even without grouping information. When grouping information is available, group IHT enhances model selection. The results displayed in Table 4 reflect $n = 1,000$ samples, $p = 10,000$ SNPs, and 100 replicates. Each SNP belongs to 1 of 500 disjoint groups containing 20 SNPs each; $j = 5$ groups are assigned $k = 3$ causal SNPs with effect sizes randomly chosen from $\{-0.2, 0.2\}$. In all there 15 causal SNPs. The indices j and k are taken as known. As described in methods, the simulated data show LD within each group, with the degree of LD between two SNPs decreasing as their separation increases.

	Ungrouped IHT	Grouped IHT
Normal	75.3 ± 13.3	80.9 ± 10.6
Logistic	25.3 ± 10.3	41.4 ± 14.8
Poisson	79.8 ± 13.0	82.6 ± 11.7
Neg Bin	75.8 ± 12.0	80.8 ± 13.4

Table 4: Percent of correlated predictors recovered ($\pm$ 1 standard error) using IHT. Doubly-sparse group IHT appears to enhance model selection.

4.6 Introduction of Prior Weights

This section considers how scaling by prior weights helps in model selection. Table 5 compares weighted IHT reconstructions with unweighted reconstructions where all weights $w_i = 1$. The weighted version of IHT consistently finds approximately 10% more true predictors than the unweighted version. Here we simulated 50 replicates involving 1000 subjects, 10,000 uncorrelated variants, and $k = 10$ true predictors for each GLM. For the sake of simplicity, we defined a prior weight $w_i = 2$ for about one-tenth of all variants, including the 10 true predictors. For the remaining SNPs the prior weight is $w_i = 1$. For instance, one can imagine that a tenth of all genotyped variants fall in a protein coding region, including the 10 true predictors, and that these variants are twice as likely to influence a trait as those falling in non-protein coding regions.

	Unweighted	Weighted
Normal	9.22 ± 0.42	9.44 ± 0.50
Logistic	7.36 ± 0.66	7.92 ± 0.53
Poisson	8.00 ± 0.61	8.24 ± 0.66
Neg Bin	9.18 ± 0.44	9.38 ± 0.49

Table 5: Comparison of unweighted and weighted IHT shows weighting finds more predictors on average ($\pm$ 1 standard error). The true number of SNPs is $k = 10$.

4.7 Data analysis of Cardiovascular Related Phenotypes from the 1966 Northern Finland Birth Cohort

We tested IHT on data from the 1966 Northern Finland Birth Cohort (NFBC1966) (32). Although this dataset is relatively modest with 5402 participants and 364,590 SNPs, it has two virtues. First, it featured in our pre-

Trait	Debias?	SNP	Position	$\hat{\beta}$	Known?
HDL	d/nd	rs9261224	30121866	-0.03	No
	d/nd	rs6917603	30125050	0.17	(24; 19)
	d/nd	rs9261256	30129920	-0.07	(19)
	d/nd	rs7120118	47242866	-0.03	(24; 32; 19)
	d/nd	rs1532085	56470658	-0.04	(24; 32; 19)
	d/nd	rs3764261	55550825	-0.05	(24; 32; 19)
	d	rs7499892	55564091	0.03	(24; 19)
	d/nd	rs3852700	65829359	-0.03	No
	nd	rs1800961	42475778	0.03	(24)
LDL	d/nd	rs646776	109620053	0.03	(32; 19; 24)
	d/nd	rs6917603	30125050	-0.05	(19; 24)

Table 6: Results generated by running IHT on high density lipoprotein (HDL) phenotype as a normal response and low density lipoprotein (LDL) as a binary response from the NFBC GWAS dataset. Cross validation was conducted without debiasing. Parameter estimates represent the output of the best size model run with both debiasing (d) and non-debiasing (nd).

vious IHT paper (19), so comparison with our earlier analysis is easy. Second, due to a population bottleneck (25), the participants' chromosomes exhibit more extensive linkage disequilibrium than is typically found in less isolated populations. Multivariate regression selection methods, including the lasso, have been criticized for their inability to deal with the dependence among predictors induced by LD. Therefore this dataset provides a realistic test case for doubly-sparse IHT.

We imputed missing genotypes with Mendel (21). Following (19), we excluded subjects with missing phenotypes, fasting subjects, and subjects on diabetes medication. We conducted quality control measures using the OpenMendel module `SnpArrays`(45). Based on these measures, we excluded SNPs with minor allele frequency ≤ 0.01 and Hardy Weinberg equilibrium p-values $\leq 10^{-5}$. As for non-genetic predictors, we included sex (the `sexOCPG` factor defined in (32)) as well as the first 2 principal components of the genotype matrix computed via `PLINK 2.0 alpha` (9). To put predictors, genetic and non-genetic, on an equal footing, we standardized all predictors to have mean zero and unit variance.

4.7.1 High Density Lipoprotein (HDL)

Using IHT we find previously associated SNPs as well as a few new potential associations. We model the HDL phenotype as normally-distributed and find a best model size $\hat{k} = 9$ based on 5-fold cross validation across model sizes $k = \{1, 2, ..., 20\}$. Without debiasing, the analysis was completed in 2 hours and 4 minutes with 30 CPU cores. Table 6 depicts the final models obtained via both debiasing (d) and non-debiasing (nd).

Importantly, IHT is able to simultaneously recover effects for SNPs (1) rs9261224, (2) rs6917603, and (3) rs6917603 with pairwise correlations of $r_{1,2} = 0.618, r_{1,3} = 0.984$, and $r_{2,3} = 0.62$. This result is achieved without grouping of SNPs, which can further increase association power. Compared with earlier analyses of these data, we find 3 SNPs that were not listed in our previous IHT paper (19), presumably due to slight algorithmic modifications. The authors of NFBC (32) found 5 SNPs associated with HDL. We did not find SNPs rs2167079 and rs255049. To date, rs255049 was replicated in a reanalysis of the same dataset (15).

SNP rs2167079 has been reported to be associated with an unrelated phenotype (30).

4.7.2 Low Density Lipoprotein (LDL) as a Binary Response

Unfortunately we did not have access to any qualitative phenotypes for this cohort, so for purposes of illustration, we fit a logistic regression model to a derived dichotomous phenotype, high versus low levels of Low Density Lipoprotein (LDL). The original data are continuous, so we choose 145 mg/dL, the midpoint between the borderline-high and high LDL cholesterol categories, to separate the two categories (16). This dichotomization resulted in 932 cases (high LDL) and 3961 controls (low LDL). Under 5-fold cross validation without debiasing across model sizes $k = \{1, 2, ..., 20\}$, we find $\hat{k} = 3$. Based 30 CPU cores, our algorithm finishes in 1 hours and 7 minutes.

Despite the loss of information inherent in dichotomization, our results are comparable to the prior results under a normal model for the original quantitative LDL phenotype. Our final model still recovers two SNP predictors with and without debiasing (Table 6). We miss all but one of the SNPs that the NFBC analysis found to be associated with LDL treated as a quantitative trait. Notably we again find an association with SNP rs6917603 that they did not report.

5 Discussion

Multivariate methods like IHT provide a principled way of model fitting and variable selection. With increasing computing power and better software, multivariate methods are likely to prevail over univariate ones. This paper introduces a scalable implementation of iterative hard thresholding for generalized linear models. Although we focused our attention on GWAS, our GLM implementation accepts arbitrary numeric data and is suitable for a variety of applied statistics problems. Our real data analysis and simulation studies suggest that IHT can (a) recover highly correlated SNPs, (b) avoid over-fitting, (c) deliver better true positive and false positive rates than either marginal testing or lasso regression, (d) recover unbiased regression coefficients, and (e) exploit prior information and group-sparsity. In our opinion, the time is ripe for the genomics community to embrace multivariate methods as a supplement to and possibly a replacement of marginal analysis.

The potential applications of iterative hard thresholding reach far beyond gene mapping. Genetics and the broader field of bioinformatics are blessed with rich, ultra-high dimensional data. IHT is designed to solve such problems. By extending IHT to the realm of generalized linear models, it becomes possible to fit regression models with more exotic distributions than the Gaussian distributions implicit in ordinary linear regression. In our view IHT will eventually join and probably supplant lasso regression as the method of choice in GWAS and other high-dimensional regression settings.

6 Methods

6.1 Data Simulation

Our simulations mimic scenarios for a range of rare and common SNPs with or without LD. Unless otherwise stated, we designate 10 SNPs to be causal with effect sizes of $0.1, 0.2, ..., 1.0$.

To generate independent SNP genotypes, we first sample a minor allele frequency $\rho_j \sim \text{Uniform}(0, 0.5)$ for each SNP j. To construct the genotype of person i at SNP j, we then sample from a binomial distribution with success probability ρ_j and two trials. The vector of genotypes (minor allele counts) for person i form row $\mathbf{x}_i^t$ of the design matrix $\mathbf{X}$. To generate SNP genotypes with linkage disequilibrium, we divide all SNPs into blocks of length 20. Within each block, we first sample $x_1 \sim \text{Bernoulli}(0.5)$. Then we form a single haplotype block of length 20 by the following Markov chain procedure:

$$x_{i+1} = \begin{cases} x_i & \text{with probability } \ p \\ 1 - x_i & \text{with probability } \ 1-p \end{cases}$$

with default $p = 0.75$. For each block we form a pool of 20 haplotypes using this procedure, ensuring every one of the 40 alleles (2 at each SNP) are represented at least once. For each person, the genotype vector in a block is formed by sampling 2 haplotypes with replacement from the pool and summing the number of minor alleles at each SNP.

Depending on the simulation, the number of subjects range from $1,000$ to $120,000$, and the number of independent SNPs range from $10,000$ to $1,000,000$. We simulate data under four GLM distributions: normal (Gaussian), Bernoulli, Poisson, and negative binomial. We generate component y_i of the response vector $\mathbf{y}$ by sampling from the corresponding distribution with mean $\mu_i = g(\mathbf{x}_i^t \boldsymbol{\beta})$, where g is the inverse link function. For normal models we assume unit variance, and for negative binomial models we assume 10 required failures. To avoid overflows, we clamp the mean $g(\mathbf{x}_i^t \boldsymbol{\beta})$ to stay within $[-20, 20]$. (See Ad Hoc Tactics for a detailed explanation). We apply the canonical link for each distribution, except for the negative binomial, where we apply the log link.

6.2 Linear Algebra with Compressed Genotype Files

The genotype count matrix stores minor allele counts. The PLINK genotype compression protocol (9) compactly stores the corresponding 0's, 1's, and 2's in 2 bits per SNP, achieving a compression ratio of 32:1 compared to storage as floating point numbers. For a sparsity level k model, we use OpenBLAS (a highly optimized linear algebra library) to compute predicted values. This requires transforming the k pertinent columns of $\mathbf{X}$ into a floating point matrix $\mathbf{X}_k$ and multiplying it times the corresponding entries $\boldsymbol{\beta}_k$ of $\boldsymbol{\beta}$. The inverse link is then applied to $\mathbf{X}_k \boldsymbol{\beta}_k$ to give the mean vector $\boldsymbol{\mu} = g(\mathbf{X}_k \boldsymbol{\beta}_k)$. In computing the GLM gradient (equation 3.2), formation of the vector $\mathbf{W}_1(\mathbf{y} - \boldsymbol{\mu})$ involves no matrix multiplications. Computation of the gradient $\mathbf{X}^t \mathbf{W}_1(\mathbf{y} - \boldsymbol{\mu})$ is more complicated because the full matrix $\mathbf{X}$ can no longer be avoided. Fortunately, the OpenMendel module `SnpArrays` can be invoked to perform compressed matrix times vector multiplication. Calculation of the steplength of IHT requires computation of the quadratic form $\nabla L(\boldsymbol{\beta}_n)^t \mathbf{X}^t \mathbf{W}_2 \mathbf{X} \nabla L(\boldsymbol{\beta}_n)$. Given the gradient, this computation requires a single compressed matrix times vector multiplication. Finally, good statistical practice calls for standardizing covariates. To standardize the genotype counts for SNP j, we estimate its minor allele frequency p_j and then substitute the ratio $\frac{x_{ij} - 2p_j}{\sqrt{2p_j(1-p_j)}}$ for the genotype count x_{ij} for person i at SNP j. This procedure is predicated on a binomial distribution for the count x_{ij}. Our previous paper (19) shows how to accommodate standardization in the matrix operations of IHT without actually forming or storing the standardized matrix.

Although multiplication via the OpenMendel module `SnpArrays` (45) is slower than OpenBLAS multiplication on small data sets, it can be as much as 10 times faster on large data sets. OpenBLAS has advantages in parallelization, but it requires floating point arrays. Once the genotype matrix $\mathbf{X}$ exceeds the memory avail-

able in RAM, expensive data swapping between RAM and hard disk memory sets in. This dramatically slows matrix multiplication. `SnpArrays` is less vulnerable to this hazard owing to compression. Once compressed data exceeds RAM, `SnpArrays` also succumbs to the swapping problem. Current laptop and desktop computers seldom have more than 32 GB of RAM, so we must resort to cluster or cloud computing when input files exceeding 32 GB.

6.3 Computations Involving Non-genetic Covariates

Non-genetic covariates are stored as double or single precision floating point entries in an $n \times r$ design matrix $\mathbf{Z}$. To accommodate an intercept, the first column should be a vector of 1's. Let γ denote the r vector of regression coefficients corresponding to $\mathbf{Z}$. The full design matrix is the block matrix $(\mathbf{X}\,\mathbf{Z})$. Matrix multiplications involving $(\mathbf{X}\,\mathbf{Z})$ should be carried out via

$$(\mathbf{X}\,\mathbf{Z})\begin{pmatrix}\boldsymbol{\beta}\\ \gamma\end{pmatrix} \;=\; \mathbf{X}\boldsymbol{\beta}+\mathbf{Z}\gamma \quad \text{and} \quad (\mathbf{X}\,\mathbf{Z})^t\,\mathbf{v} \;=\; \begin{pmatrix}\mathbf{X}^t\mathbf{v}\\ \mathbf{Z}^t\mathbf{v}\end{pmatrix}.$$

Adherence to these rules ensures a low memory footprint. Multiplication involving $\mathbf{X}$ can be conducted as previously explained. Multiplication involving $\mathbf{Z}$ can revert to BLAS.

6.4 Parallel Computation

The OpenBLAS library accessed by Julia is inherently parallel. Beyond that we incorporate parallel processing in cross validation. Recall that in q-fold cross validation we separate subjects into q disjoint subsets. We then fit a training model using $q-1$ of those subsets and record the mean-squared prediction error on the omitted subset. Each of the q subsets serve as the testing set exactly once. Testing error is averaged across the different folds for each sparsity levels k. The lowest average testing error determines the recommended sparsity. Cross validation offers opportunities for parallelism because the different folds can be handled on different processors. `MendelIHT.jl` uses one of Julia's (4) standard library `Distributed.jl` to load different folds to different CPUs.

6.5 Ad Hoc Tactics to Prevent Overflows

In Poisson and negative binomial regressions, the inverse link argument $\exp(\mathbf{x}_i^t\boldsymbol{\beta})$ experiences numerical overflows when the inner product $\mathbf{x}_i^t\boldsymbol{\beta}$ is too large. In general, we avoid running Poisson regression when response means are large. In this regime a normal approximation is preferred. As a safety feature, `MendelIHT.jl` clamps values of $\mathbf{x}_i^t\boldsymbol{\beta}$ to the interval $[-20, 20]$. Note that penalized regression suffers from the same overflow catastrophes.

6.6 Convergence and Backtracking

For each proposed IHT step we check whether the objective $L(\boldsymbol{\beta})$ increases. When it does not, we step-halve at most 5 times to restore the ascent property. Convergence is declared when

$$\frac{||\boldsymbol{\beta}_{n+1} - \boldsymbol{\beta}_n||_\infty}{||\boldsymbol{\beta}_n||_\infty + 1} < \text{Tolerance},$$

with the default tolerance being 0.0001. The addition of 1 in the denominator of the convergence criterion guards against division by 0.

7 Availability of source code

Project name: MendelIHT
Project home page:
https://github.com/biona001/MendelIHT.jl
Operating systems: Mac OS, Linux
Programming language: Julia 1.0, 1.1
License: MIT

8 Availability of supporting data and materials

The Northern Finland Birth Cohort 1966 (NFBC1966) (32) was downloaded from dbGaP under dataset accession pht002005.v1.p1. The code to generate simulated data, as well as their subsequent analysis, are available in our github repository under *figures* folder.

9 Declarations

9.1 List of abbreviations

GWAS: genome wide association studies; SNP: single nucleotide polymorphism; IHT: iterative hard threhsolding; GLM: generalized linear models; LD: linkage disequilibrium; MAF: minor allele frequency; Neg Bin: negative binomial; NFBC: northern finland birth cohort; HDL: high density lipoprotein; LDL: low density lipoprotein;

9.2 Ethics, Consent for publication, competing interest

The authors declare no conflicts of interest. As described in (32), informed consent from all study subjects was obtained using protocols approved by the Ethical Committee of the Northern Ostrobothnia Hospital District.

9.3 Funding

Benjamin Chu was supported by NIH T32-HG002536 training grant and the 2018 Google Summer of Code. Kevin Keys was supported by a diversity supplement to NHLBI grant R01HL135156, the UCSF Bakar Computational Health Sciences Institute, the Gordon and Betty More Foundation grant GBMF3834, and the Alfred P. Sloan Foundation grant 2013-10-27 to UC Berkeley through the Moore-Sloan Data Sciences Environment initiative at the Berkeley Institute for Data Science (BIDS). Kenneth Lange was supported by grants from the National Human Genome Research Institute (HG006139) and the National Institute of General Medical Sciences (GM053275). Janet Sinsheimer was supported by grants from the National Institute of General Medical Sciences (GM053275), the National Human Genome Research Institute (HG009120) and the National Science Foundation (DMS-1264153).

9.4 Author's Contributions

JSS, KK, KL, BC contributed to the design of the study, interpretation of results, and writing of the manuscript. BC designed and implemented the simulations and conducted the data analyses. BC and KK developed the software. KL and BC developed the algorithms.

10 Acknowledgements

We thank UCLA for providing the computing resources via the Hoffman2 cluster.

The NFBC1966 Study is conducted and supported by the National Heart, Lung, and Blood Institute (NHLBI) in collaboration with the Broad Institute, UCLA, University of Oulu, and the National Institute for Health and Welfare in Finland. This manuscript was not prepared in collaboration with investigators of the NFBC1966 Study and does not necessarily reflect the opinions or views of the NFBC1966 Study Investigators, Broad Institute, UCLA, University of Oulu, National Institute for Health and Welfare in Finland and the NHLBI.

References

[1] David H. Alexander and Kenneth Lange. Stability selection for genome-wide association. *Genetic Epidemiology*, 35:722–728, 2011.

[2] Amir Beck. *Introduction to nonlinear optimization: Theory, algorithms, and applications with MATLAB*, volume 19. Siam, 2014.

[3] Amir Beck and Marc Teboulle. A linearly convergent algorithm for solving a class of nonconvex/affine feasibility problems. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pages 33–48. Springer, 2011.

[4] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59:65–98, 2017.

[5] Thomas Blumensath and Mike E Davies. Normalized iterative hard thresholding: Guaranteed stability and performance. *IEEE Journal of Selected Topics in Signal Processing*, 4:298–309, 2010.

[6] Patrick Breheny and Jian Huang. Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Annals of Applied Statistics*, 5:232–253, 2011.

[7] William S Bush and Jason H Moore. Genome-wide association studies. *PLoS Computational Biology*, 8:e1002822, 2012.

[8] Rita M Cantor, Kenneth Lange, and Janet S Sinsheimer. Prioritizing GWAS results: a review of statistical methods and recommendations for their application. *The American Journal of Human Genetics*, 86:6–22, 2010.

[9] Christopher C Chang, Carson C Chow, Laurent CAM Tellier, Shashaank Vattikuti, Shaun M Purcell, and James J Lee. Second-generation PLINK: rising to the challenge of larger and richer datasets. *GigaScience*, 4:7, 2015.

[10] de Lamare, Rodrigo C. Knowledge-aided normalized iterative hard thresholding algorithms and applications to sparse reconstruction. *arXiv preprint arXiv:1809.09281*, 2018.

[11] Annette J Dobson and Adrian Barnett. *An introduction to generalized linear models*. Chapman and Hall/CRC, 2008.

[12] Simon Foucart. Hard thresholding pursuit: an algorithm for compressive sensing. *SIAM Journal on Numerical Analysis*, 49:2543–2563, 2011.

[13] Jerome Friedman, Trevor Hastie, and Rob Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33:1, 2010.

[14] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*, 2010.

[15] Lisa Gai and Eleazar Eskin. Finding associated variants in genome-wide association studies on multiple traits. *Bioinformatics*, 34:i467–i474, 2018.

[16] Scott M Grundy, Neil J Stone, Alison L Bailey, Craig Beam, Kim K Birtcher, Roger S Blumenthal, Lynne T Braun, Sarah de Ferranti, Joseph Faiella-Tommasino, Daniel E Forman, and et al. 2018 AHA/ACC/AACVPR/AAPA/ABC/ACPM/ADA/AGS/APhA/ASPC/NLA/PCNA guideline on the management of blood cholesterol: a report of the American College of Cardiology/American Heart Association Task Force on Clinical Practice Guidelines. *Journal of the American College of Cardiology*, 73:e285–e350, 2019.

[17] Buhm Han and Eleazar Eskin. Random-effects model aimed at discovering associations in meta-analysis of genome-wide association studies. *The American Journal of Human Genetics*, 88:586–598, 2011.

[18] Gabriel E. Hoffman, Benjamin A. Logsdon, and Jason G. Mezey. PUMA: A unified framework for penalized multiple regression analysis of GWAS data. *PLoS Computational Biology*, 9:e1003101, 06 2013.

[19] Kevin L Keys, Gary K Chen, and Kenneth Lange. Iterative hard thresholding for model selection in genome-wide association studies. *Genetic Epidemiology*, 41:756–768, 2017.

[20] Kenneth Lange. *Numerical analysis for statisticians*. Springer Science & Business Media, 2010.

[21] Kenneth Lange, Jeanette C Papp, Janet S Sinsheimer, Ram Sripracha, Hua Zhou, and Eric M Sobel. Mendel: the swiss army knife of genetic analysis programs. *Bioinformatics*, 29:1568–1570, 2013.

[22] Dahua Lin, John Myles White, Simon Byrne, Douglas Bates, Andreas Noack, John Pearson, Alex Arslan, Kevin Squire, David Anthoff, Theodore Papamarkou, Mathieu Besançon, and et al. JuliaStats/Distributions.jl: a Julia package for probability distributions and associated functions, may 2019.

[23] Po-Ru Loh, George Tucker, Brendan K Bulik-Sullivan, Bjarni J Vilhjalmsson, Hilary K Finucane, Rany M Salem, Daniel I Chasman, Paul M Ridker, Benjamin M Neale, Bonnie Berger, Nick Patterson, and Alkes L Price. Efficient bayesian mixed-model analysis increases association power in large cohorts. *Nature Genetics*, 47:284, 2015.

[24] Jacqueline MacArthur, Emily Bowler, Maria Cerezo, Laurent Gil, Peggy Hall, Emma Hastings, Heather Junkins, Aoife McMahon, Annalisa Milano, Joannella Morales, and et al. The new NHGRI-EBI Catalog of published genome-wide association studies (GWAS Catalog). *Nucleic acids research*, 45:D896–D901, 2016.

[25] Alicia R Martin, Konrad J Karczewski, Sini Kerminen, Mitja I Kurki, Antti-Pekka Sarin, Mykyta Artomov, Johan G Eriksson, Tõnu Esko, Giulio Genovese, Aki S Havulinna, and et al. Haplotype sharing provides insights into fine-scale population history and disease in finland. *The American Journal of Human Genetics*, 102:760–775, 2018.

[26] Rahul Mazumder, Jerome H. Friedman, and Trevor Hastie. SparseNet: Coordinate descent with non-convex penalties. *Journal of the American Statistical Association*, 106:1125–1138, 2011.

[27] Peter McCullagh. *Generalized Linear Models*. Routledge, 2018.

[28] Lukas Meier, Sara Van De Geer, and Peter Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70:53–71, 2008.

[29] Nicolai Meinshausen and Peter Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72:417–473, 2010.

[30] Stacey Melquist, David W Craig, Matthew J Huentelman, Richard Crook, John V Pearson, Matt Baker, Victoria L Zismann, Jennifer Gass, Jennifer Adamson, Szabolcs Szelinger, et al. Identification of a novel risk locus for progressive supranuclear palsy by a pooled genomewide scan of 500,288 single-nucleotide polymorphisms. *The American Journal of Human Genetics*, 80:769–778, 2007.

[31] Deanna Needell and Joel A Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26:301–321, 2009.

[32] Chiara Sabatti, Susan K Service, Anna-Liisa Hartikainen, Anneli Pouta, Samuli Ripatti, Jae Brodsky, Chris G Jones, Noah A Zaitlen, Teppo Varilo, Marika Kaakinen, and et al. Genome-wide association analysis of metabolic traits in a birth cohort from a founder population. *Nature genetics*, 41:35, 2009.

[33] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[34] Shashaank Vattikuti, James J Lee, Christopher C Chang, Stephen DH Hsu, and Carson C Chow. Applying compressed sensing to genome-wide association studies. *GigaScience*, 3:10, 2014.

[35] Peter M Visscher, Naomi R Wray, Qian Zhang, Pamela Sklar, Mark I McCarthy, Matthew A Brown, and Jian Yang. 10 years of GWAS discovery: biology, function, and translation. *The American Journal of Human Genetics*, 101:5–22, 2017.

[36] Tong Tong Wu, Yi Fang Chen, Trevor Hastie, Eric Sobel, and Kenneth Lange. Genome-wide association analysis by lasso penalized logistic regression. *Bioinformatics*, 25:714–721, 2009.

[37] Tong Tong Wu and Kenneth Lange. Coordinate descent algorithms for lasso penalized regression. *The Annals of Applied Statistics*, 2:224–244, 2008.

[38] Jason Xu, Eric Chi, and Kenneth Lange. Generalized Linear Model Regression under Distance-to-set Penalties. In *Advances in Neural Information Processing Systems 30.*, pages 1385–1395. Curran Associates, Inc., 2017.

[39] Fan Yang, Rina Foygel Barber, Prateek Jain, and John Lafferty. Selective inference for group-sparse linear models. In *Advances in Neural Information Processing Systems*, pages 2469–2477, 2016.

[40] Xiao-Tong Yuan, Ping Li, and Tong Zhang. Gradient hard thresholding pursuit. *Journal of Machine Learning Research*, 18:166–1, 2017.

[41] Achim Zeileis, Christian Kleiber, and Simon Jackman. Regression models for count data in R. *Journal of Statistical Software*, 27:1–25, 2008.

[42] Tong Zhang. Analysis of multi-stage convex relaxation for sparse regularization. *Journal of Machine Learning Research*, 11:1081–1107, 2010.

[43] Hua Zhou, David H Alexander, Mary E Sehl, Janet S Sinsheimer, Eric M Sobel, and Kenneth Lange. Penalized regression for genome-wide association screening of sequence data. In *Pacific Symposium on Biocomputing*, pages 106–117. World Scientific, 2011.

[44] Hua Zhou, Mary E Sehl, Janet S Sinsheimer, and Kenneth Lange. Association screening of common and rare genetic variants by penalized regression. *Bioinformatics*, 26:2375, 2010.

[45] Hua Zhou, Janet S Sinsheimer, Douglas M Bates, Benjamin B Chu, Christopher A German, Sarah S Ji, Kevin L Keys, Juhyun Kim, Seyoon Ko, Gordon D Mosher, and et al. OpenMendel: a cooperative programming project for statistical genetics. *Human Genetics*, pages 1–11, 2019.