
Variance function estimation in high-dimensions

Mladen Kolar
James Sharpnack

MLADENK@CS.CMU.EDU
JSHARPNA@CS.CMU.EDU

Machine Learning Department, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213 USA

Abstract

We consider the high-dimensional heteroscedastic regression model, where the mean and the log variance are modeled as a linear combination of input variables. Existing literature on high-dimensional linear regression models has largely ignored non-constant error variances, even though they commonly occur in a variety of applications ranging from biostatistics to finance. In this paper we study a class of non-convex penalized pseudolikelihood estimators for both the mean and variance parameters. We show that the Heteroscedastic Iterative Penalized Pseudolikelihood Optimizer (HIPPO) achieves the oracle property, that is, we prove that the rates of convergence are the same as if the true model was known. We demonstrate numerical properties of the procedure on a simulation study and real world data.

1. Introduction

High-dimensional regression models have been studied extensively in both machine learning and statistical literature. Statistical inference in high-dimensions, where the sample size n is smaller than the ambient dimension p , is impossible without assumptions. As the concept of parsimony is important in many scientific domains, most of the research in the area of high-dimensional statistical inference is done under the assumption that the underlying model is sparse, in the sense that the number of relevant parameters is much smaller than p , or that it can be well approximated by a sparse model.

MK is partially supported through the grants NIH R01GM087694 and AFOSR FA9550010247. JS is partially supported by AFOSR under grant FA95501010382.

Appearing in *Proceedings of the 29th International Conference on Machine Learning*, Edinburgh, Scotland, UK, 2012. Copyright 2012 by the author(s)/owner(s).

Penalization of the empirical loss by the ℓ_1 norm has become a popular tool for obtaining sparse models and huge amount of literature exists on theoretical properties of estimation procedures (see, e.g., Zhao & Yu, 2006; Wainwright, 2009; Zhang, 2009; Zhang & Huang, 2008, and references therein) and on efficient algorithms that numerically find estimates (see Bach et al., 2011, for an extensive literature review). Due to limitations of the ℓ_1 norm penalization, high-dimensional inference methods based on the class of concave penalties have been proposed that have better theoretical and numerical properties (see, e.g., Fan & Li, 2001; Fan & Lv, 2009; Lv & Fan, 2009; Zhang & Zhang, 2011).

In all of the above cited work, the main focus is on model selection and mean parameter estimation. Only few papers deal with estimation of the variance in high-dimensions (Sun & Zhang, 2011; Fan et al., 2012) although it is a fundamental problem in statistics. Variance appears in the confidence bounds on estimated regression coefficients and is important for variable selection as it appears in Akaike's information criterion (AIC) and the Bayesian information criterion (BIC). Furthermore, it provides confidence on the predictive performance of a forecaster.

In applied regression it is often the case that the error variance is non-constant. Although the assumption of a constant variance can sometimes be achieved by transforming the dependent variable, e.g., by using a Box-Cox transformation, in many cases transformation does not produce a constant error variance (Carroll & Ruppert, 1988). Another approach is to ignore the heterogeneous variance and use standard estimation techniques, but such estimators are less efficient. Aside from its use in reweighting schemes, estimating variance is important because the resulting prediction intervals become more accurate and it is often important to explore which input variables drive the variance. In this paper, we will model the variance directly as a parametric function of the explanatory variables.

Heteroscedastic regression models are used in a variety of fields ranging from biostatistics to economet-

rics, finance and quality control in manufacturing. In this paper, we study penalized estimation in high-dimensional heteroscedastic linear regression models, where the mean and the log variance are modeled as a linear combination of explanatory variables. Modeling the log variance as a linear combination of the explanatory variables is a common choice as it guarantees positivity and is also capable of capturing variance that may vary over several orders of magnitudes (Carroll & Ruppert, 1988; Harvey, 1976). Main contributions of this paper are as follows. First, we propose HIPPO (Heteroscedastic Iterative Penalized Pseudolikelihood Optimizer) for estimation of both the mean and variance parameters. Second, we establish the oracle property (in the sense of Fan & Lv (2009)) for the estimated mean and variance parameters. Finally, we demonstrate numerical properties of the proposed procedure on a simulation study, where it is shown that HIPPO outperforms other methods, and analyze a real data set.

1.1. Problem Setup and Notation

Consider the usual heteroscedastic linear model,

$$y_i = \mathbf{x}'_i \boldsymbol{\beta} + \sigma(\mathbf{x}_i, \boldsymbol{\theta}) \epsilon_i, \quad i = 1, \dots, n, \quad (1)$$

where $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)' = (\mathbf{X}_1, \dots, \mathbf{X}_p)$ is an $n \times p$ matrix of predictors with *i.i.d.* rows $\mathbf{x}_1, \dots, \mathbf{x}_n$, $\mathbf{y} = (y_1, \dots, y_n)$ is an n -vector of responses, the vectors $\boldsymbol{\beta} \in \mathbb{R}^p$ and $\boldsymbol{\theta} \in \mathbb{R}^p$ are p -vectors of mean and variance parameters, respectively, and $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)$ is an n -vector of *i.i.d.* random noise with mean 0 and variance 1. We assume that the noise $\boldsymbol{\epsilon}$ is independent of the predictors \mathbf{X} . The function $\sigma(\mathbf{x}, \boldsymbol{\theta})$ has a known parametric form and, for simplicity of presentation, we assume that it takes a particular form $\sigma(\mathbf{x}_i, \boldsymbol{\theta}) = \exp(\mathbf{x}'_i \boldsymbol{\theta} / 2)$.

Throughout the paper we use $[n]$ to denote the set $\{1, \dots, n\}$. For any index set $S \subseteq [p]$, we denote $\boldsymbol{\beta}_S$ to be the subvector containing the components of the vector $\boldsymbol{\beta}$ indexed by the set S , and \mathbf{X}_S denotes the submatrix containing the columns of \mathbf{X} indexed by S . For a vector $\mathbf{a} \in \mathbb{R}^n$, we denote $\text{supp}(\mathbf{a}) = \{j : a_j \neq 0\}$ the support set, $\|\mathbf{a}\|_q$, $q \in (0, \infty)$, the ℓ_q -norm defined as $\|\mathbf{a}\|_q = (\sum_{i \in [n]} a_i^q)^{1/q}$ with the usual extensions for $q \in \{0, \infty\}$, that is, $\|\mathbf{a}\|_0 = |\text{supp}(\mathbf{a})|$ and $\|\mathbf{a}\|_\infty = \max_{i \in [n]} |a_i|$. For notational simplicity, we denote $\|\cdot\| = \|\cdot\|_2$ the ℓ_2 norm. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times p}$ we denote $\|\mathbf{A}\|_2$ the operator norm, $\|\mathbf{A}\|_F$ the Frobenius norm, and $\Lambda_{\min}(\mathbf{A})$ and $\Lambda_{\max}(\mathbf{A})$ denote the smallest and largest eigenvalue respectively.

Under the model in (1), we are interested in estimating both $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$. In high-dimensions, when $p \gg n$, it is

common to assume that the support $\boldsymbol{\beta}$ is small, that is, $S = \text{supp}(\boldsymbol{\beta})$ and $|S| \ll n$. Similarly, we assume that the support $T = \text{supp}(\boldsymbol{\theta})$ is small.

1.2. Related Work

Consider the model (1) with constant variance, i.e., $\sigma(\mathbf{x}, \boldsymbol{\theta}) \equiv \sigma_0$. Most of the existing high-dimensional literature is focused on estimation of the mean parameter $\boldsymbol{\beta}$ in this homoscedastic regression model. Under a variety of assumptions and regularity conditions, any penalized estimation procedure mentioned in introduction can, in theory, select the correct sparse model with probability tending to 1. Literature on variance estimation is not as developed. Fan et al. (2012) proposed a two step procedure for estimation of the unknown variance σ_0 , while (Sun & Zhang, 2011) proposed an estimation procedure that jointly estimates the model and the variance.

Problem of estimation in the heteroscedastic linear regression models have been studied extensively in the classical setting with p fixed, however, the problem of estimation under the model (1) when $p \gg n$ has not been adequately studied. Jia et al. (2010) assume that $\sigma(\mathbf{x}, \boldsymbol{\theta}) = |\mathbf{x}' \boldsymbol{\beta}|$ and show that Lasso is sign consistent for the mean parameter $\boldsymbol{\beta}$ under certain conditions. Their study shows limitations of lasso, for which many highly scalable solvers exist. However, no new methodology is developed, as the authors acknowledge that the log-likelihood function is highly non-convex. Dette & Wagener (2011) study the adaptive lasso under the model in (1). Under certain regularity conditions, they show that the adaptive lasso is consistent, with sub-optimal asymptotic variance. However, the weighted adaptive lasso is both consistent and achieves optimal asymptotic variance, under the assumption that the variance function is consistently estimated. However, they do not discuss how to obtain an estimator of the variance function in a principled way and resort to an ad-hoc fitting of the residuals. Daye et al. (2011) develop HHR procedure that optimizes the penalized log-likelihood under (1) with the ℓ_1 -norm penalty on both the mean and variance parameters. As the objective is not convex, HHR estimates $\boldsymbol{\beta}$ with $\boldsymbol{\theta}$ fixed and then estimates $\boldsymbol{\theta}$ with $\boldsymbol{\beta}$ fixed, until convergence. Since the objective is biconvex, HHR converges to a stationary point. However, no theory is provided for the final estimates.

2. Methodology

In this paper, we propose HIPPO (Heteroscedastic Iterative Penalized Pseudolikelihood Optimizer) for estimating $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ under model (1).

In the first step, HIPPO finds the penalized pseudolikelihood maximizer of β by solving the following objective

$$\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} \|\mathbf{y} - \mathbf{X}\beta\|^2 + 2n \sum_{j \in [p]} \rho_{\lambda_S}(|\beta_j|), \quad (2)$$

where ρ_{λ_S} is the penalty function and the tuning parameter λ_S controls the sparsity of the solution $\hat{\beta}$.

In the second step, HIPPO forms the penalized pseudolikelihood estimate for θ by solving

$$\hat{\theta} = \arg \min_{\theta \in \mathbb{R}^p} \sum_{i \in [n]} \mathbf{x}'_i \theta + \sum_{i \in [n]} \hat{\eta}_i^2 \exp(-\mathbf{x}'_i \theta) + 4n \sum_{j \in [p]} \rho_{\lambda_T}(|\theta_j|) \quad (3)$$

where $\hat{\eta} = \mathbf{y} - \mathbf{X}\hat{\beta}$ is the vector of residuals.

Finally, HIPPO computes the reweighted estimator of the mean by solving

$$\hat{\beta}_w = \arg \min_{\beta \in \mathbb{R}^p} \sum_{i \in [n]} \frac{(y_i - \mathbf{x}'_i \beta)^2}{\hat{\sigma}_i} + 2n \sum_{j \in [p]} \rho_{\lambda_S}(|\beta_j|) \quad (4)$$

where $\hat{\sigma}_i = \exp(\mathbf{x}'_i \hat{\theta}/2)$ are the weights.

In classical literature, estimation under heteroscedastic models is achieved by employing a pseudolikelihood objective. The pseudolikelihood maximization principle prescribes the scientist to maximize a surrogate likelihood, i.e. one that is believed to be similar to the likelihood with the true unknown fixed variances (or means alternatively). In classical theory, central limit theorems are derived for many pseudo-maximum likelihood (PML) estimators using generalized estimating equations (Ziegler, 2011). HIPPO fits neatly into the pseudolikelihood framework because the first step is a regularized PML where only the mean structure needs to be correctly specified. The second step and third steps may be similarly cast as PML estimators. Indeed, all our theoretical results are due to the fact that in each step we are optimizing a pseudolikelihood that is similar to the true unknown likelihoods (with alternating free parameters). Moreover, it is known that if the surrogate variances in the mean PML are more similar to the true variances then the resulting estimates will be more asymptotically efficient. With this in mind, we recommend a third reweighting procedure with the variance estimates from the second step.

Fan & Li (2001) advocate usage of penalty functions that result in estimates satisfying three properties: unbiasedness, sparsity and continuity. A reasonable estimator should correctly identify the support of the true

parameter with probability converging to one. Furthermore, on this support, the estimated coefficients should have the same asymptotic distribution as if an estimator that knew the true support was used. Such an estimator satisfies the *oracle property*. A number of concave penalties result in estimates that satisfy this property: the SCAD penalty (Fan & Li, 2001), the MCP penalty (Zhang, 2010) and a class of folded concave penalties (Lv & Fan, 2009). For concreteness, we choose to use the SCAD penalty, which is defined by its derivative

$$\rho'_\lambda(t) = \lambda \left[I\{t \leq \lambda\} + \frac{(a\lambda - t)_+}{(a-1)\lambda} I\{t > \lambda\} \right], \quad (5)$$

where often $a = 3.7$ is used. Note that estimates produced by the ℓ_1 -norm penalty are biased, and hence this penalty does not achieve oracle property.

HIPPO is related to the iterative HHR algorithm of Daye et al. (2011). In particular, the first two iterations of HHR are equivalent to HIPPO with the SCAD penalty replaced with the ℓ_1 norm penalty. In practice, one can continue iterating between solving (3) and (4), however, establishing theoretical properties for those iterates is a non-trivial task. From our numerical studies, we observe that HIPPO performs well when stopped after the first two iterations.

2.1. Tuning Parameter Selection

As described in the previous section, HIPPO requires selection of the tuning parameters λ_S and λ_T , which balance the complexity of the estimated model and the fit to data. A common approach is to form a grid of candidate values for the tuning parameters λ_S and λ_T and chose those that minimize the AIC or BIC criterion

$$\text{AIC}(\lambda_S, \lambda_T) = \sum_{i \in [n]} \ell(y_i, \mathbf{x}_i; \hat{\beta}, \hat{\theta}) + 2\hat{d}f, \quad (6)$$

$$\text{BIC}(\lambda_S, \lambda_T) = \sum_{i \in [n]} \ell(y_i, \mathbf{x}_i; \hat{\beta}, \hat{\theta}) + \hat{d}f \log n \quad (7)$$

where, up to constants,

$$\ell(y, \mathbf{x}; \beta, \theta) = \mathbf{x}'\theta + (y - \mathbf{x}'\beta)^2 \exp(-\mathbf{x}'\theta)$$

is the negative log-likelihood and

$$\hat{d}f = |\text{supp}(\hat{\beta})| + |\text{supp}(\hat{\theta})|$$

is the estimated degrees of freedom. In Section 4, we compare performance of the AIC and the BIC for HIPPO in a simulation study.

2.2. Optimization Procedure

In this section, we describe numerical procedures used to solve optimization problems in (2), (3) and (4). Our procedures are based on the local linear approximation for the SCAD penalty developed in (Zou & Li, 2008), which gives:

$$\rho_\lambda(|\beta_j|) \approx \rho_\lambda(|\beta_j^{(k)}|) + \rho'_\lambda(|\beta_j^{(k)}|)(|\beta_j| - |\beta_j^{(k)}|),$$

for $\beta_j \approx \beta_j^{(k)}$.

This approximation allows us to substitute the SCAD penalty $\sum_{j \in [p]} \rho_\lambda(|\beta_j|)$ in (2), (3) and (4) with

$$\sum_{j \in [p]} \rho'_\lambda(|\hat{\beta}_j^{(k)}|)|\beta_j|, \quad (8)$$

and iteratively solve each objective until convergence of $\{\hat{\beta}^{(k)}\}_k$. We set the initial estimates $\hat{\beta}^{(0)}$ and $\hat{\theta}^{(0)}$ to be the solutions of the ℓ_1 -norm penalized problems. The convergence of these iterative approximations follows from the convergence of the MM (minorize-maximize) algorithms (Zou & Li, 2008).

With the approximation of the SCAD penalty given in (8), we can solve (2) and (4) using standard lasso solvers, e.g., we use the proximal method of Beck & Teboulle (2009). The objective in (3) is minimized using a coordinate descent algorithm, which is detailed in Daye et al. (2011).

3. Theoretical Properties of HIPPO

In this section, we present theoretical properties of HIPPO. In particular, we show that HIPPO achieves the oracle property for estimating the mean and variance under the model (1). All the proofs are deferred to Appendix.

We will analyze HIPPO under the following assumptions, which are standard in the literature on high-dimensional statistical learning (see, e.g. Fan et al., 2012).

Assumption 1. The matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)' \in \mathbb{R}^{n \times p}$ has independent rows that satisfy $\mathbf{x}_i = \boldsymbol{\Sigma}^{1/2} \mathbf{z}_i$ where $\{\mathbf{z}_i\}_i$ are *i.i.d.* subgaussian random variables with $\mathbb{E} \mathbf{z}_i = \mathbf{0}$, $\mathbb{E} \mathbf{z}_i \mathbf{z}_i' = \mathbf{I}$ and parameter K (see Appendix for more details on subgaussian random variables). Furthermore, there exist two constants $C_{\min}, C_{\max} > 0$ such that

$$0 < C_{\min} \leq \Lambda_{\min}(\boldsymbol{\Sigma}) \leq \Lambda_{\max}(\boldsymbol{\Sigma}) \leq C_{\max} < \infty.$$

Assumption 2. The errors $\epsilon_1, \dots, \epsilon_n$ are *i.i.d.* subgaussian with zero mean and parameter 1.

Assumption 3. There are two constants $\bar{\beta}$ and $\bar{\theta}$ such that $\|\boldsymbol{\beta}\| \leq \bar{\beta} < \infty$ and $\|\boldsymbol{\theta}\| \leq \bar{\theta} < \infty$.

Assumption 4. $|S| = C_S n^{\alpha_S}$ and $|T| = C_T n^{\alpha_T}$ for some $\alpha_S \in (0, 1)$ and $\alpha_T \in (0, 1/3)$ and constants $C_S, C_T > 0$.

The following assumption will be needed for showing the consistency of the weighted estimator $\hat{\boldsymbol{\beta}}_w$ in (4).

Assumption 5. Define

$$\mathbf{D}_{SS} = n^{-1} \mathbf{X}'_S \text{diag}(\exp(-\mathbf{X}\boldsymbol{\theta})) \mathbf{X}_S.$$

There exist constants $0 \leq D_{\min}, D_{\max} \leq \infty$ such that

$$\lim_{n \rightarrow \infty} \mathbb{P}[\Lambda_{\max}(\mathbf{D}_{SS}) \leq D_{\max}] = 1, \quad \text{and}$$

$$\lim_{n \rightarrow \infty} \mathbb{P}[\Lambda_{\min}(\mathbf{D}_{SS}) \geq D_{\min}] = 1.$$

Furthermore, we have that

$$\lim_{n \rightarrow \infty} \|\mathbf{D}_{SS} - \mathbb{E} \mathbf{D}_{SS}\|_2 = o_P(1).$$

With these assumption, we state our first result, regarding the estimator $\hat{\boldsymbol{\beta}}$ in (2).

Theorem 1. *Suppose that the assumptions (1)-(4) are satisfied. Furthermore, assume that $\lambda_S \geq c_1 \sqrt{\log(p) \exp(\sqrt{c_2 \log(n)})/n}$, $\min_{j \in [S]} |\beta_j| \gg \lambda_S \gg c_3 \sqrt{\log(s) \exp(\sqrt{c_2 \log(n)})/n}$ and $\log(p) = \mathcal{O}(n^{\alpha_0})$ for some $\alpha_0 \in (0, 1)$. Then there is a strict local minimizer $\hat{\boldsymbol{\beta}} = (\hat{\boldsymbol{\beta}}'_S, \mathbf{0}'_{S^c})'$ of (2) that satisfies*

$$\|\hat{\boldsymbol{\beta}}_S - \boldsymbol{\beta}_S\|_\infty \leq c_3 \sqrt{\frac{\exp(\sqrt{c_2 \log(n)}) \log(s)}{n}} \quad (9)$$

for some positive constants c_1, c_2 , and c_3 and sufficiently large n .

In addition, if we suppose that assumption (5) is satisfied, then for any fixed $\mathbf{a} \in \mathbb{R}^s$ with $\|\mathbf{a}\|_2 = 1$ the following weak convergence holds

$$\frac{\sqrt{n}}{\zeta} \mathbf{a}' (\hat{\boldsymbol{\beta}}_S - \boldsymbol{\beta}_S) \xrightarrow{D} \mathcal{N}(0, 1) \quad (10)$$

where $\zeta^2 = \mathbf{a}' \boldsymbol{\Sigma}_{SS}^{-1} \mathbb{E} \mathbf{D}_{SS} \boldsymbol{\Sigma}_{SS}^{-1} \mathbf{a}$.

The first result stated in Theorem 1 established that $\hat{\boldsymbol{\beta}}$ achieves the weak oracle property in the sense of (Lv & Fan, 2009). The extra term $\exp(\sqrt{\log n})$ is subpolynomial in n and appears in the bound (9) due to the heteroscedastic nature of the errors. The second result establishes the strong oracle property of the estimator $\hat{\boldsymbol{\beta}}$ in the sense of (Fan & Lv, 2009), that is, we establish the asymptotic normality on the true support

S . The asymptotic normality shows that $\hat{\beta}_S$ has the same asymptotic variance as the ordinary least squares (OLS) estimator on the true support. However, in the case of a heteroscedastic model the OLS estimator is dominated by the generalized least squares estimator. Later in this section, we will demonstrate that $\hat{\beta}_w$ has better asymptotic variance. Note that $\hat{\beta}$ correctly selects the mean model and estimates the parameters at the correct rate. From the upper and lower bounds on λ_S , we see how the rate at which p can grow and the minimum coefficient size are related. Larger the ambient dimension p gets, larger the size of λ_S , which lower bounds the size of the minimum coefficient.

Our next result establishes correct model selection for the variance parameter θ .

Theorem 2. *Suppose that assumptions (1)-(5) are satisfied. Suppose further that $\lambda_T \geq n^{\alpha_T-1/2} \log(p) \log(n)$ and $\min_{j \in [T]} |\theta_j| \geq \lambda_T$. Then there is a strict local minimizer $\hat{\theta} = (\hat{\theta}'_T, \mathbf{0}'_{T^c})'$ with the strong oracle property,*

$$\|n^{(1-\alpha_T)/2}(\hat{\theta} - \theta)\| = O_{\mathbb{P}}(1)$$

Moreover, for any fixed $\mathbf{a} \in \mathbb{R}^t$ with $\|\mathbf{a}\|_2 = 1$ the following weak convergence holds

$$\frac{\sqrt{n}}{\zeta} \mathbf{a}'(\hat{\theta}_T - \theta_T) \xrightarrow{D} \mathcal{N}(0, 1) \quad (11)$$

where $\zeta^2 = \mathbf{a}' \Sigma_{TT}^{-1} \mathbf{a}$.

With the convergence result of $\hat{\theta}$ we can prove consistency and asymptotic normality of the weighted estimator $\hat{\beta}$ in (4).

Theorem 3. *Suppose that the assumptions (1)-(5) are satisfied and that there exists an estimator $\hat{\theta}$ satisfying $\|\hat{\theta} - \theta\|_2 = \mathcal{O}(r_n)$, for a sequence $r_n \rightarrow 0$ and $\text{supp}(\hat{\theta}) = \text{supp}(\theta)$. Furthermore, assume that $\lambda_S \geq c_1 \sqrt{\log(p) \exp(\sqrt{c_2 \log(n)})/n}$, $\min_{j \in [S]} |\beta_j| \gg \lambda_S \gg c_3 r_n \exp(\sqrt{c_2 \log(n)}) \log(n)$ and $\log(p) = \mathcal{O}(n^{\alpha_0})$ for some $\alpha_0 \in (0, 1)$. Then there is a strict local minimizer $\hat{\beta}_w = (\hat{\beta}'_{w,S}, \mathbf{0}'_{S^c})$ of (4) that satisfies*

$$\|\hat{\beta}_{w,S} - \beta_S\|_{\infty} \leq c_3 r_n \exp(\sqrt{c_2 \log(n)}) \log(n) \quad (12)$$

for some positive constants c_1, c_2 , and c_3 and sufficiently large n .

Furthermore, for any fixed $\mathbf{a} \in \mathbb{R}^s$ with $\|\mathbf{a}\|_2 = 1$ the following weak convergence holds

$$\frac{\sqrt{n}}{\zeta_w} \mathbf{a}'(\hat{\beta} - \beta) \xrightarrow{D} \mathcal{N}(0, 1) \quad (13)$$

where $\zeta_w^2 = \mathbf{a}'(\mathbb{E}D_{SS})^{-1} \mathbf{a}$.

	$\ \theta - \hat{\theta}\ _2$	Pre $_{\theta}$	Rec $_{\theta}$
$\rho = 0$			
HHR-AIC	0.59(0.13)	0.4(0.17)	1.00(0.00)
HIPPO-AIC	0.26(0.15)	0.6(0.22)	1.00(0.00)
HHR-BIC	0.59(0.13)	0.39(0.16)	1.00(0.00)
HIPPO-BIC	0.26(0.15)	0.59(0.22)	1.00(0.00)
$\rho = 0.5$			
HHR-AIC	0.32(0.12)	0.68(0.21)	1.00(0.00)
HIPPO-AIC	0.38(0.22)	0.69(0.25)	1.00(0.03)
HHR-BIC	0.32(0.12)	0.68(0.21)	1.00(0.00)
HIPPO-BIC	0.38(0.22)	0.69(0.25)	0.99(0.03)

Table 1. Mean (sd) performance of HHR and HIPPO under the model in Example 1 (averaged over 100 independent runs). The mean parameter β is assumed to be known.

Theorem 3 establishes convergence of the weighted estimator $\hat{\beta}_w$ in (4) and the model selection consistency. The rate of convergence depends on the rate of convergence of the variance estimator, r_n . From Theorem 2, we show the parametric rate of convergence for $\hat{\theta}_S$. The second result of Theorem 3 states that the weighted estimator $\hat{\beta}_{w,S}$ is asymptotically normal, with the same asymptotic variance as the generalized least squares estimator which knows the true model and variance function $\sigma(\mathbf{x}, \theta)$.

4. Monte-Carlo Simulations

In this section, we conduct two small scale simulation studies to demonstrate finite sample performance of HIPPO. We compare it to the HHR procedure (Daye et al., 2011).

Convergence of the parameters is measured in the ℓ_2 norm, $\|\hat{\beta} - \beta\|$ and $\|\hat{\theta} - \theta\|$. We measure the identification of the support of β and θ using precision and recall. Let \hat{S} denote the estimated set of non-zero coefficients of S , then the precision is calculated as $\text{Pre}_{\beta} := |\hat{S} \cap S|/|\hat{S}|$ and the recall as $\text{Rec}_{\beta} := |\hat{S} \cap S|/|S|$. Similarly, we can define precision and recall for the variance coefficients. We report results averaged over 100 independent runs.

4.1. Example 1

Assume that the data is generated iid from the following model $Y = \sigma(\mathbf{X})\epsilon$ where ϵ follows a standard normal distribution and the logarithm of the variance is given by

$$\log \sigma(\mathbf{X})^2 = X_1 + X_2 + X_3.$$

Variance function estimation in high-dimensions

	#it	$\ \beta - \hat{\beta}\ _2$	Pre_β	Rec_β	$\ \theta - \hat{\theta}\ _2$	Pre_θ	Rec_θ
$n = 200$							
HHR-AIC	1st	0.78(0.52)	0.44(0.22)	1.00(0.00)	2.10(0.11)	0.25(0.10)	0.54(0.16)
	2nd	0.31(0.13)	0.88(0.15)	1.00(0.00)	1.80(0.16)	0.29(0.07)	0.71(0.14)
HIPPO-AIC	1st	0.66(0.84)	0.75(0.29)	1.00(0.02)	2.00(0.16)	0.20(0.10)	0.52(0.16)
	2nd	0.08(0.07)	0.84(0.24)	1.00(0.00)	1.50(0.30)	0.30(0.11)	0.75(0.12)
HHR-BIC	1st	0.77(0.48)	0.58(0.17)	1.00(0.00)	2.10(0.10)	0.41(0.18)	0.45(0.14)
	2nd	0.31(0.13)	0.89(0.13)	1.00(0.00)	1.90(0.16)	0.38(0.15)	0.65(0.17)
HIPPO-BIC	1st	0.70(0.83)	0.80(0.25)	0.99(0.03)	2.00(0.14)	0.39(0.18)	0.50(0.17)
	2nd	0.08(0.06)	0.97(0.07)	1.00(0.00)	1.60(0.28)	0.44(0.16)	0.72(0.14)
$n = 400$							
HHR-AIC	1st	0.59(0.37)	0.58(0.26)	1.00(0.00)	1.90(0.11)	0.36(0.14)	0.72(0.18)
	2nd	0.30(0.24)	0.98(0.06)	1.00(0.00)	1.70(0.16)	0.43(0.13)	0.81(0.16)
HIPPO-AIC	1st	0.44(0.54)	0.87(0.22)	1.00(0.00)	1.80(0.18)	0.28(0.10)	0.67(0.15)
	2nd	0.06(0.29)	0.97(0.12)	1.00(0.02)	1.00(0.31)	0.56(0.18)	0.93(0.09)
HHR-BIC	1st	0.59(0.37)	0.66(0.20)	1.00(0.00)	1.90(0.11)	0.46(0.18)	0.66(0.20)
	2nd	0.30(0.23)	0.98(0.06)	1.00(0.00)	1.70(0.17)	0.46(0.13)	0.80(0.17)
HIPPO-BIC	1st	0.46(0.58)	0.89(0.19)	1.00(0.01)	1.80(0.18)	0.39(0.17)	0.65(0.17)
	2nd	0.06(0.29)	0.99(0.06)	1.00(0.02)	1.00(0.31)	0.63(0.20)	0.92(0.09)

Table 2. Mean (sd) performance of HHR and HIPPO under the model in Example 2 (averaged over 100 independent runs). We report estimated models after the first and second iteration.

The covariates associated with the variance are jointly normal with equal correlation ρ , and marginally $\mathcal{N}(0,1)$. The remaining covariates, X_4, \dots, X_p are iid random variables following the standard Normal distribution and are independent from (X_1, X_2, X_3) . We set $(n, p) = (200, 2000)$ and use $\rho = 0$ and $\rho = 0.5$. Estimation procedures know that $\beta = \mathbf{0}$ and we only estimate the variance parameter θ . This example is provided to illustrate performance of the penalized pseudolikelihood estimators in an idealized situation. When the mean parameter needs to be estimated as well, we expect the performance of the procedures only to get worse. Since the mean is known, both HHR and HIPPO only solve the optimization procedure in (3), HHR with the ℓ_1 -norm penalty and HIPPO with the SCAD penalty, without iterating between (4) and (3).

Table 1 summarizes the results. Under this toy model, we observe that HIPPO performs better than HHR when the correlation between the relevant predictors is $\rho = 0$. However, we do not observe the difference between the two procedures when $\rho = 0.5$. The difference between the AIC and BIC is already visible in this example when $\rho = 0$. The AIC tends to pick more complex models, while the BIC is more conservative and selects a model with fewer variables.

4.2. Example 2

The following non-trivial model is borrowed from Daye et al. (2011). The response variable Y satisfies

$$Y = \beta_0 + \sum_{j \in [p]} X_j \beta_j + \exp(\theta_0 + \sum_{j \in [p]} X_j \theta_j) \epsilon$$

with $p = 600$, $\beta_0 = 2$, $\theta_0 = 1$,

$$\beta_{[12]} = (3, 3, 3, 1.5, 1.5, 1.5, 0, 0, 0, 2, 2, 2)'$$

$$\theta_{[15]} = (1, 1, 1, 0, 0, 0, 0.5, 0.5, 0.5, 0, 0, 0, 0.75, 0.75, 0.75)'$$

and the remainder of the coefficients are 0. The covariates are jointly Normal with $\text{cov}(X_i, X_j) = 0.5^{|i-j|}$ and the error ϵ follows the standard Normal distribution. This is a more realistic model than the one described in the previous example. We set $p = 600$ and the number of samples $n = 200$ and $n = 400$.

Table 2 summarizes results of the simulation. We observe that HIPPO consistently outperforms HHR in all scenarios. Again, a general observation is that the AIC selects more complex models although the difference is less pronounced when the sample size $n = 400$. Furthermore, we note that the estimation error significantly reduces after the first iteration, which demonstrates final sample benefits of estimating the variance.

Recall that Theorem 1 proves that the estimate $\hat{\beta}$ consistently estimates the true parameter β . However, it is important to estimate the variance parameter θ well, both in theory (see Theorem 3) and practice.

5. Real Data Application

Forecasting the gross domestic product (GDP) of a country based on macroeconomic indicators is of significant interest to the economic community. We obtain both the country GDP figures (specifically we use the GDP per capita using current prices in units of a ‘national currency’) and macroeconomic variables from the International Monetary Fund’s World Economic Outlook (WEO) database. The WEO database contains records for macroeconomic variables from 1980 to 2016 (with forecasts).

To form our response variable, $\mathbf{y}_{i,t}$, we form log-returns of the GDP for each country (i) for each time point (t) after records began and before the forecasting commenced (each country had a different year at which forecasting began). After removing missing values, we obtained 31 variables that can be grouped into a few broad categories: balance of payments, government finance and debt, inflation, and demographics. We apply various transformations, including lagging and logarithms forming the vectors $\mathbf{x}_{i,t}$. We fit the heteroscedastic AR(1) model with HIPPO.

$$\mathbf{y}_{i,t} = \mathbf{x}'_{i,t-1}\beta + \exp(\mathbf{x}'_{i,t-1}\theta)\epsilon_{i,t}$$

In order to initially assess the heteroscedasticity of the data, we form the LASSO estimator with the LARS package in R selecting with BIC. It is common practice when diagnosing heteroscedasticity to plot the studentized residuals against the fitted values. We bin the bulk of the samples into three groups by fitted values, and observe the box-plot of each bin by residuals (Figure 2). It is apparent that there is a difference of variances between these bins, which is corroborated by performing a F-test of equal variances across the second and third bins (p-value of 4×10^{-6}). We further observe differences of variance between country GDP log returns. We analyzed the distribution of responses separated by countries: Canada, Finland, Greece and Italy. The p-value from the F-test for equality of variances between the countries Canada and Greece is 0.008, which is below even the pairwise Bonferroni correction of 0.0083 at 0.05 significance level. This demonstrates heteroscedasticity in the WEO dataset, and we are justified in fitting non-constant variance.

We compare the results from HIPPO and HHR when applied to the WEO data set. The tuning parameters were selected with BIC over a grid for λ_S and

	HIPPO	HHR
MSE	0.0089	0.0091
$-\ell(\mathbf{y}, \mathbf{X}; \hat{\beta}, \hat{\theta})$	0.4953	0.6783
$ \hat{S} $	5.4	8.9
$ \hat{T} $	8.2	5.1

Table 3. Performance of HIPPO and HHR on WEO data averaged over 10 folds.

λ_T . The metrics used to compare the algorithms are mean square error (MSE) defined by $\frac{1}{n} \sum_i (y_{i,t} - \hat{y}_{i,t})^2$, the partial prediction score defined as the average of the negative log likelihoods, and the number of selected mean parameters and variance parameters. We perform 10-fold cross validation to obtain unbiased estimates of these metrics. In Table 3 we observe that HIPPO outperforms HHR in terms of MSE and partial prediction score.

6. Discussion

We have addressed the problem of statistical inference in high-dimensional linear regression models with heteroscedastic errors. Heteroscedastic errors arise in many applications and industrial settings, including biostatistics, finance and quality control in manufacturing. We have proposed HIPPO for model selection and estimation of both the mean and variance parameters under a heteroscedastic model. HIPPO can be deployed naturally into an existing data analysis workflow. Specifically, as a first step, a statistician performs penalized estimation of the mean parameters and then, as a second step, tests for heteroscedasticity by running the second step of HIPPO. If heteroscedasticity is discovered, HIPPO can then be used to solve penalized generalized least squares objective. Furthermore, HIPPO is well motivated from the penalized pseudolikelihood maximization perspective and achieves the oracle property in high-dimensional problems.

Throughout the paper, we focus on a specific parametric form of the variance function for simplicity of presentation. Our method can be extended to any parametric form, however, the assumptions will become more cumbersome and the particular numerical procedure would change. It is of interest to develop general unified framework for estimation of arbitrary parametric form of the variance function. Another open research direction includes non-parametric estimation of the variance function in high-dimensions, which could be achieved with sparse additive models (see Ravikumar et al., 2009).

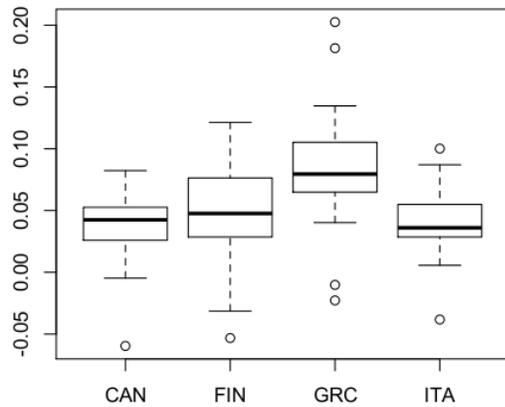


Figure 1. A box-plot of the GDP log-returns for the 4 countries with the most observed time points (Canada, Finland, Greece, and Italy).

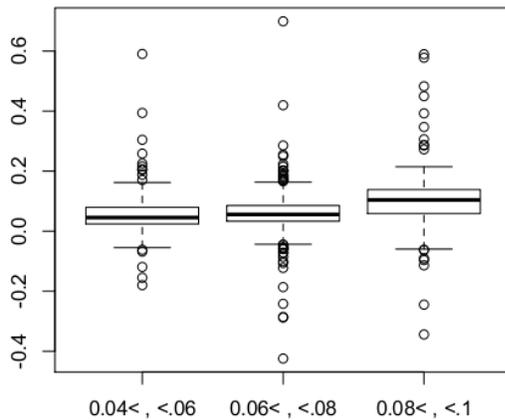


Figure 2. A box-plot of the studentized residuals binned by LASSO predicted $y_{i,t}$. Only the segment of the predicted response with the bulk of the samples was binned; the breaks in the bins are at 0.04, 0.06, 0.08, and 0.1.

References

Bach, F., Jenatton, R., Mairal, J., and Obozinski, G. Optimization with Sparsity-Inducing Penalties. *ArXiv e-prints*, 2011.

Beck, A. and Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183202, 2009.

Carroll, R.J. and Ruppert, D. *Transformation and weighting in regression*, volume 30. Chapman & Hall/CRC, 1988.

Daye, Z. John, Chen, Jinbo, and Li, Hongzhe. High-dimensional heteroscedastic regression with an application to eqtl data analysis. *Biometrics*, pp. no-no, 2011. ISSN 1541-0420.

Detle, H. and Wagener, J. The adaptive lasso in high

dimensional sparse heteroscedastic models. 2011.

Fan, J. and Li, R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96:1348–1360, December 2001.

Fan, J. and Lv, J. Non-Concave Penalized Likelihood with NP-Dimensionality. *ArXiv e-prints*, October 2009.

Fan, Jianqing, Guo, Shaojun, and Hao, Ning. Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(1):37–65, 2012. ISSN 1467-9868.

Harvey, A. C. Estimating regression models with multiplicative heteroscedasticity. *Econometrica*, 44(3):461–466, 1976.

Jia, J., Rohe, K., and Yu, B. The lasso under heteroscedasticity. *Arxiv preprint arXiv:1011.1026*, 2010.

Lv, Jinchi and Fan, Yingying. A unified approach to model selection and sparse recovery using regularized least squares. *Annals of Statistics*, 37(6A):3498–3528, 2009.

Ravikumar, P., Lafferty, J., Liu, H., and Wasserman, L. Sparse additive models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(5):1009–1030, 2009.

Sun, T. and Zhang, C.-H. Scaled Sparse Linear Regression. *ArXiv e-prints*, April 2011.

Wainwright, Martin J. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso). *IEEE Transactions on Information Theory*, 55(5):2183–2202, May 2009.

Zhang, C.-H. and Zhang, T. A General Theory of Concave Regularization for High Dimensional Sparse Estimation Problems. *ArXiv e-prints*, 2011.

Zhang, C.H. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2): 894–942, 2010.

Zhang, Cun-Hui and Huang, Jian. The sparsity and bias of the lasso selection in high-dimensional linear regression. *Annals of Statistics*, 36(4):1567–1594, 2008.

Zhang, T. Some sharp performance bounds for least squares regression with L1 regularization. *Annals of Statistics*, 37(5A):2109–2144, 2009.

Zhao, Peng and Yu, Bin. On model selection consistency of lasso. *J. Mach. Learn. Res.*, 7:2541–2563, 2006.

Ziegler, Andreas. *Generalized Estimating Equations*. Number 204 in Lecture Notes in Statistics. Springer, 2011.

Zou, Hui and Li, Runze. One-step sparse estimates in non-concave penalized likelihood models. *Annals of Statistics*, 36(4):1509–1533, 2008.