

STRUCTURED VARIABLE SELECTION AND ESTIMATION

BY MING YUAN^{*}, V. ROSHAN[†] JOSEPH AND HUI ZOU[‡]

Georgia Institute of Technology and University of Minnesota

In linear regression problems with related predictors, it is desirable to do variable selection and estimation by maintaining the hierarchical or structural relationships among predictors. In this paper, we propose nonnegative garrote methods that can naturally incorporate such relationships defined through effect heredity principles or marginality principles. We show that the methods are very easy to compute and enjoy nice theoretical properties. We also show that the methods can be easily extended to deal with more general regression problems such as generalized linear models. Simulations and real examples are used to illustrate the merits of the proposed methods.

1. Introduction. When considering regression with a large number of predictors, variable selection becomes important. Numerous methods have been proposed in the literature for the purpose of variable selection, ranging from the classical information criteria such as AIC and BIC to regularization based modern techniques such as the nonnegative garrote (Breiman, 1995), the Lasso (Tibshirani, 1996) and the SCAD (Fan and Li, 2001) among many others. Although these methods enjoy excellent performance in many appli-

^{*}The research of Ming Yuan was supported in part by NSF grants DMS-MPSA-0624841 and DMS-0706724.

[†]The research of V. Roshan Joseph was supported in part by NSF grant CMMI-0448774.

[‡]The research of Hui Zou was supported in part by NSF grant DMS-0706733.

Keywords and phrases: Effect heredity, nonnegative garrote, quadratic programming, regularization, variable selection.

cations, they do not take the hierarchical or structural relationship among predictors into account and therefore can lead to models that are hard to interpret.

Consider, for example, multiple linear regression with both main effects and two-way interactions where a dependent variable Y and q explanatory variables X_1, X_2, \dots, X_q are related through

$$(1.1) \quad Y = \beta_1 X_1 + \dots + \beta_q X_q + \beta_{11} X_1^2 + \beta_{12} X_1 X_2 + \dots + \beta_{qq} X_q^2 + \epsilon,$$

where $\epsilon \sim \mathcal{N}(0, \sigma^2)$. Commonly used general purpose variable selection techniques, including those mentioned above, do not distinguish interactions $X_i X_j$ from main effects X_i and can select a model with an interaction but neither of its main effects, i.e., $\beta_{ij} \neq 0$ and $\beta_i = \beta_j = 0$. It is therefore useful to invoke the so-called effect heredity principle (Hamada and Wu 1992) in this situation. There are two popular versions of the heredity principle (Chipman 1996). Under *strong heredity*, for a two-factor interaction effect $X_i X_j$ to be active both its parent effects, X_i and X_j , should be active; whereas under *weak heredity* only one of its parent effects needs to be active. Likewise, one may also require that X_i^2 can be active only if X_i is also active. The strong heredity principle is closely related to the notion of marginality (Nelder, 1977; McCullagh and Nelder, 1989; Nelder, 1994) which ensures that the response surface is invariant under scaling and translation of the explanatory variables in the model. Interested readers are also referred to McCullagh (2002) for a rigorous discussion about what criteria a sensible statistical model should obey. Li, Sudarsanam and Frey (2006) recently

conducted a meta-analysis of 113 data sets from published factorial experiments and concluded that an overwhelming majority of these real studies conform with the heredity principles. This clearly shows the importance of using these principles in practice.

These two heredity concepts can be extended to describe more general hierarchical structure among predictors. With slight abuse of notation, write a general multiple linear regression as

$$(1.2) \quad Y = X\beta + \epsilon,$$

where $X = (X_1, X_2, \dots, X_p)$, and $\beta = (\beta_1, \dots, \beta_p)'$. Throughout this paper, we center each variable so that the observed mean is zero and therefore, the regression equation has no intercept. In its most general form, the hierarchical relationship among predictors can be represented by sets $\{\mathcal{D}_i : i = 1, \dots, p\}$, where \mathcal{D}_i contains the parent effects of the i th predictor. For example, the dependence set of $X_i X_j$ is $\{X_i, X_j\}$ in the quadratic model (1.1). In order that the i th variable can be considered for inclusion, all elements of \mathcal{D}_i must be included under strong heredity principle, and at least one element of \mathcal{D}_i should be included under weak heredity principle. Other types of heredity principle such as the partial heredity principle (Nelder, 1998) can also be incorporated in this framework. The readers are referred to Yuan, Joseph and Lin (2007) for further details. As pointed out by Turlach (2004), it could be very challenging to conform with the hierarchical structure in the popular variable selection methods. In this paper, we specifically address this issue and consider how to effectively impose such hierarchical structures among

the predictors in variable selection and coefficient estimation, which we refer to as *structured variable selection and estimation*.

Despite its great practical importance, structured variable selection and estimation has received only scant attention in the literature. Earlier interests in structured variable selection come from the analysis of designed experiments where heredity principles have proven to be powerful tools in resolving complex aliasing patterns. Hamada and Wu (1992) introduced a modified stepwise variable selection procedure that can enforce effect heredity principles. Later, Chipman (1996) and Chipman, Hamada and Wu (1997) discussed how the effect heredity can be accommodated in the stochastic search variable selection method developed by George and McCulloch (1993). See also Joseph and Delaney (2007) for another Bayesian approach. Despite its elegance, the Bayesian approach can be computationally demanding for large scale problems. Recently, Yuan, Joseph and Lin (2007) proposed generalized LARS algorithms (Osborne, Presnell and Turlach, 2000; Efron et al., 2004) to incorporate heredity principles into model selection. Efron et al. (2004) and Turlach (2004) also considered alternative strategies to enforce the strong heredity principle in the LARS algorithm. Compared with earlier proposals, the generalized LARS procedures enjoy tremendous computational advantages, which make them particularly suitable for problems of moderate or large dimensions. However, Yuan and Lin (2007) recently showed that LARS may not be consistent in variable selection. Moreover, the generalized LARS approach is not flexible enough to incorporate many

of the hierarchical structures among predictors. More recently, Zhao, Rocha and Yu (2006) and Choi, Li and Zhu (2006) proposed penalization methods to enforce the strong heredity principle in fitting a linear regression model. However, it is not clear how to generalize them to handle more general heredity structures and their theoretical properties remain unknown.

In this paper, we propose a new framework for structured variable selection and estimation that complements and improves over the existing approaches. We introduce a family of shrinkage estimator that is similar in spirit to the nonnegative garrote, which Yuan and Lin (2007) recently showed to enjoy great computational advantages, nice asymptotic properties and excellent finite sample performance. We propose to incorporate structural relationships among predictors as linear inequality constraints on the corresponding shrinkage factors. The resulting estimates can be obtained as the solution of a quadratic program and very efficiently solved using the standard quadratic programming techniques. We show that unlike LARS, it is consistent in both variable selection and estimation provided that the true model has such structures. Moreover, the linear inequality constraints can be easily modified to adapt to any situation arising in practical problems and therefore, is much more flexible than the existing approaches. We also extend the original nonnegative garrote as well as the proposed structured variable selection and estimation methods to deal with the generalized linear models.

The proposed approach is much more flexible than the generalized LARS

approach in Yuan et al. (2007). For example, suppose a group of variables is expected to follow strong heredity and another group weak heredity, then in the proposed approach we only need to use the corresponding constraints for strong and weak heredity in solving the quadratic program. Whereas the approach of Yuan et al. (2007) is algorithmic and therefore, require considerable amount of expertise with the generalized LARS code to implement these special needs. However, there is a price to be paid for this added flexibility; it is not as fast as the generalized LARS.

The rest of the paper is organized as follows. We introduce the methodology and study its asymptotic properties in the next section. In Section 3, we extend the methodology to generalized linear models. Section 4 discusses the computational issues involved in the estimation. Simulations and real data examples are presented in Sections 5 and 6 to illustrate the proposed methods. We conclude with some discussions in Section 7.

2. Structured Variable Selection and Estimation. The original nonnegative garrote estimator introduced by Breiman (1995) is a scaled version of the least square estimate. Given n independent copies $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$ of (X, Y) where X is a p -dimensional covariate and Y is a response variable, the shrinkage factor $\theta(M) = (\theta_1(M), \dots, \theta_p(M))'$ is given as the minimizer to

$$(2.1) \quad \frac{1}{2} \|Y - Z\theta\|^2, \quad \text{subject to} \quad \sum_{j=1}^p \theta_j \leq M, \text{ and } \theta_j \geq 0, \forall j,$$

where with slight abuse of notation, $Y = (y_1, \dots, y_n)'$, $Z = (\mathbf{z}_1, \dots, \mathbf{z}_n)'$, and \mathbf{z}_i is a p dimensional vector whose j th element is $x_{ij}\widehat{\beta}_j^{\text{LS}}$ and $\widehat{\beta}^{\text{LS}}$ is the least square estimate based on (1.2). Here $M \geq 0$ is a tuning parameter. The nonnegative garrote estimate of the regression coefficient is subsequently defined as $\widehat{\beta}_j^{\text{NG}}(M) = \theta_j(M)\widehat{\beta}_j^{\text{LS}}$, $j = 1, \dots, p$. With an appropriately chosen tuning parameter M , some of the scaling factors could be estimated by exact zero and therefore the corresponding predictors are eliminated from the selected model.

2.1. *Strong Heredity Principles.* Following Yuan, Joseph and Lin (2007), let \mathcal{D}_i contain the parent effects of the i th predictor. Under the strong heredity principle, we need to impose the constraint that $\widehat{\beta}_j \neq 0$ for any $j \in \mathcal{D}_i$ if $\widehat{\beta}_i \neq 0$. A naive approach to incorporating effect heredity is therefore to minimize (2.1) under this additional constraint. However, in doing so, we lose the convexity of the optimization problem and generally will end up with problems such as multiple local optima and potentially NP hardness. Recall that the nonnegative garrote estimate of β_i is $\widehat{\beta}_i^{\text{LS}}\theta_i(M)$. Since $\widehat{\beta}_i^{\text{LS}} \neq 0$ with probability one, X_i will be selected if and only if scaling factor $\theta_i > 0$, in which case, θ_i behaves more or less like an indicator of the inclusion of X_i in the selected model. Therefore, the strong heredity principles can be enforced by requiring

$$(2.2) \quad \theta_i \leq \theta_j, \quad \forall j \in \mathcal{D}_i.$$

Note that if $\theta_i > 0$, (2.2) will force the scaling factor for all its parents to be positive and consequently active. Since these constraints are linear in terms of the scaling factor, minimizing (2.1) under (2.2) remains a quadratic program. Figure 1 illustrates the feasible region of the nonnegative garrote with such constraints in contrast with the original nonnegative garrote where no heredity rules are enforced. We consider two effects and their interaction with the corresponding shrinking factors denoted by θ_1 , θ_2 and θ_{12} respectively. In both situations, the feasible region is a convex polyhedron in the three dimensional space.

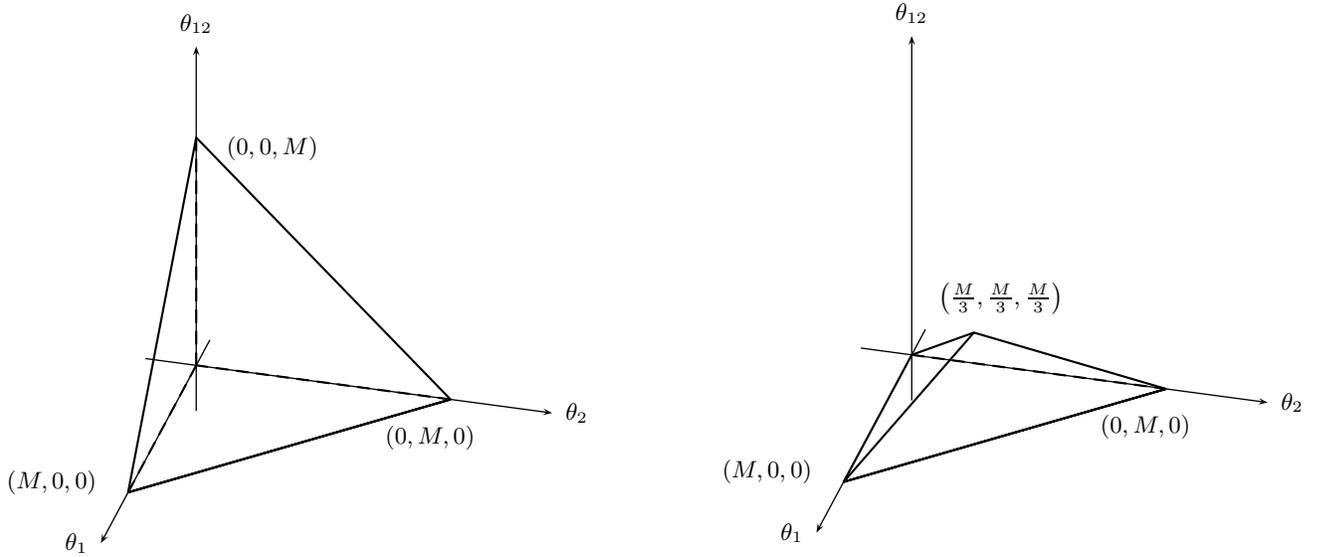


FIG 1. Feasible region of the nonnegative garrote with (right) and without (left) strong heredity constraints

2.2. *Weak Heredity Principles.* Similarly, when considering weak heredity principles, we can require that

$$(2.3) \quad \theta_i \leq \max_{j \in \mathcal{D}_i} \theta_j.$$

However the feasible region under such constraints is no longer convex as demonstrated in the left panel of Figure 2. Subsequently, minimizing (2.1) subject to (2.3) is not feasible. To overcome this problem, we suggest to use the convex envelop of these constraints for the weak heredity principles:

$$(2.4) \quad \theta_i \leq \sum_{j \in \mathcal{D}_i} \theta_j.$$

Again, these constraints are linear in terms of the scaling factor and minimizing (2.1) under (2.4) remains a quadratic program. Note that $\theta_i > 0$ implies that $\sum_{j \in \mathcal{D}_i} \theta_j > 0$ and therefore (2.4) will require at least one of its parents to be included in the model. In other words, constraint(2.4) can be employed in place of (2.3) to enforce the weak heredity principle. The small difference between the feasible regions of (2.3) and (2.4) also suggests that the selected model may only differ slightly between the two constraints. We opt for (2.4) because of the great computational advantage it brings about.

2.3. *Asymptotic Properties.* To gain further insight of the proposed structured variable selection and estimation methods, we study their asymptotic properties. We show here that the proposed methods estimate the zero coefficients by zero with probability tending to one, and at the same time give root- n consistent estimate to the nonzero coefficients provided that the true

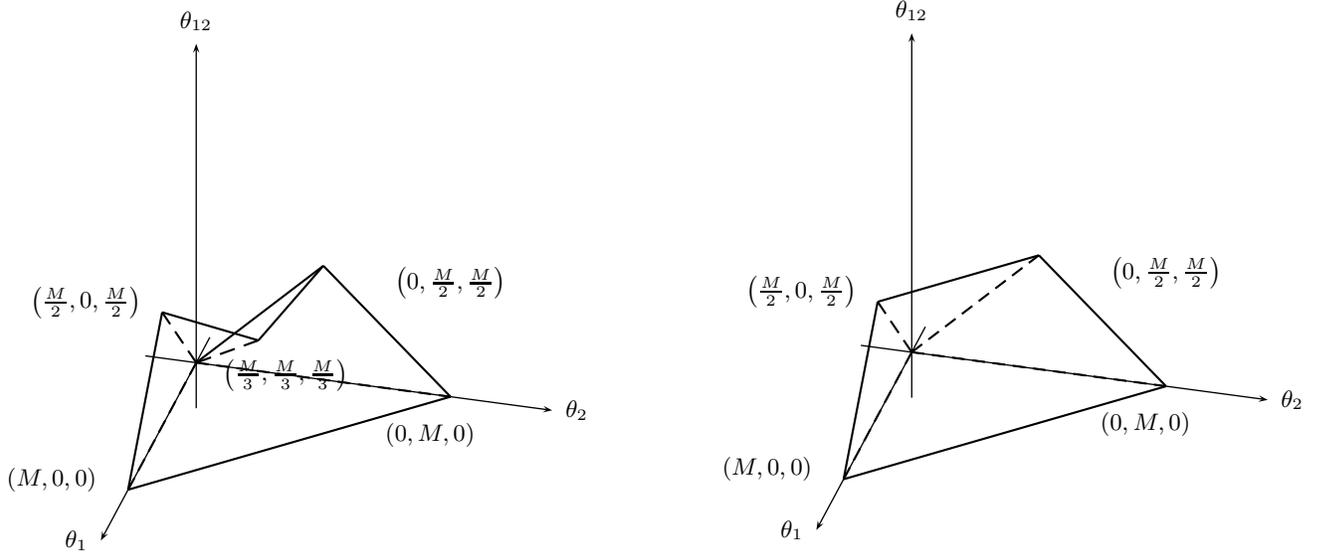


FIG 2. Feasible region of the nonnegative garrote with constraint $\theta_{12} \leq \max(\theta_1, \theta_2)$ (left) and the relaxed constraint $\theta_{12} \leq \theta_1 + \theta_2$ (right).

data generating mechanism satisfies such heredity principles. Denote by \mathcal{A} the indices of the predictors in the true model, i.e., $\mathcal{A} = \{j : \beta_j \neq 0\}$. Write $\hat{\beta}^{\text{SVS}}$ the estimate obtained from the proposed structured variable selection procedure.

Under strong heredity, the shrinkage factors $\hat{\theta}^{\text{SVS}}(M)$ can be equivalently written in the Lagrange form (Boyd and Vandenberghe, 2004)

$$\arg \min_{\theta} \left(\|Y - Z\theta\|^2 + \lambda_n \sum_{j=1}^p \theta_j \right)$$

subject to $\theta_j \geq 0, \theta_j \leq \min_{k \in \mathcal{D}_j} \theta_k$ for some Lagrange parameter $\lambda_n \geq 0$.

For the weak heredity principle, we replace the constraints $\theta_j \leq \min_{k \in \mathcal{D}_j} \theta_k$ with $\theta_j \leq \sum_{k \in \mathcal{D}_j} \theta_k$.

THEOREM 2.1. *Assume that $X'X/n \rightarrow \Sigma$ and Σ is positive definite. If the true model satisfies the strong/weak heredity principles, and $\lambda_n \rightarrow \infty$ in a fashion such that $\lambda_n = o(\sqrt{n})$ as n goes to $+\infty$, then the structured estimate with the corresponding heredity principle satisfies $P(\widehat{\beta}_j^{\text{SVS}} = 0) \rightarrow 1$ for any $j \notin \mathcal{A}$, and $\widehat{\beta}_j^{\text{SVS}} - \beta_j = O_p(n^{-1/2})$ if $j \in \mathcal{A}$.*

All the proofs can be accessed as the supplement materials. Note that when $\lambda_n = 0$, there is no penalty and the proposed estimates reduce to the least squares estimate which is consistent in estimation. The theorems suggest that if instead the tuning parameter λ_n escapes to infinity at a rate slower than \sqrt{n} , the resulting estimates not only achieve root-n consistency in terms of estimation but also are consistent in variable selection whereas the ordinary least squares estimator does not possess such model selection ability.

3. Generalized Regression. The nonnegative garrote was originally introduced for variable selection in multiple linear regression. But the idea can be extended to more general regression settings where Y depends on X through a scalar parameter $\eta(X) = \beta_0 + X\beta$ where $(\beta_0, \beta)'$ is a $(p+1)$ -dimensional unknown coefficient vector. It is worth pointing out that such extensions have not been proposed in literature so far.

A common approach to estimating η is by means of the maximum likelihood. Let $\ell(Y, \eta(X))$ be negative log conditional likelihood function of $Y|X$.

The maximum likelihood estimate is given as the minimizer of:

$$L(Y, \eta(X)) = \sum_{i=1}^n \ell(y_i, \eta(\mathbf{x}_i))$$

For example, in logistic regression,

$$\ell(y_i, \eta(\mathbf{x}_i)) = y_i \eta(\mathbf{x}_i) - \log(1 + e^{\eta(\mathbf{x}_i)}).$$

More generally, ℓ can be replaced with any loss functions such that its expectation $E(\ell(Y, \eta(X)))$ with respect to the joint distribution of (X, Y) is minimized at $\eta(\cdot)$.

To perform variable selection, we propose the following extension of the original nonnegative garrote. We use the maximum likelihood estimate $\widehat{\beta}^{\text{MLE}}$ as a preliminary estimate of β . Similar to the original nonnegative garrote, define $Z_j = X_j \widehat{\beta}_j^{\text{MLE}}$. Next we estimate the shrinkage factors by

$$(3.1) \quad (\widehat{\theta}_0, \widehat{\theta}^{\text{SVS}})' = \arg \min_{\theta_0, \theta} L(Y, Z\theta + \theta_0),$$

subject to $\sum \theta_j \leq M$ and $\theta_j \geq 0$ for any $j = 1, \dots, p$. In the case of normal linear regression, L becomes the least squares and it is not hard to see that the solution of (3.1) always satisfies $\widehat{\theta}_0 = 0$ because all variables are centered. Therefore, without loss of generality we could assume that there is no intercept in the normal linear regression. The same, however, is not true for more general L and therefore, θ_0 is included in (3.1). Our final estimate of β_j is then given as $\widehat{\beta}_j^{\text{MLE}} \widehat{\theta}_j^{\text{SVS}}(M)$ for $j = 1, \dots, p$. To impose the strong or weak heredity principle, we add additional constraints $\theta_j \leq \min_{k \in \mathcal{D}_j} \theta_k$ or $\theta_j \leq \sum_{k \in \mathcal{D}_j} \theta_k$ respectively.

Theorem 1 can also be extended to more general regression settings. Similar to before, under strong heredity,

$$\widehat{\theta}^{\text{SVS}}(M) = \arg \min_{\theta_0, \theta} \left(L(Y, Z\theta + \theta_0) + \lambda_n \sum_{j=1}^p \theta_j \right)$$

subject to $\theta_j \geq 0, \theta_j \leq \min_{k \in \mathcal{D}_j} \theta_k$ for some $\lambda_n \geq 0$. Under weak heredity principles, we use the constraints $\theta_j \leq \sum_{k \in \mathcal{D}_j} \theta_k$ instead of $\theta_j \leq \min_{k \in \mathcal{D}_j} \theta_k$.

We shall assume that the following regularity conditions hold:

- (A.1) $\ell(\cdot, \cdot)$ is a strictly convex function of the second argument;
- (A.2) The maximum likelihood estimate $\widehat{\beta}^{\text{MLE}}$ is root- n consistent;
- (A.3) The observed information matrix converges to a positive definite matrix, i.e.,

$$(3.2) \quad \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \ell''(y_i, \mathbf{x}_i \widehat{\beta}^{\text{MLE}} + \widehat{\beta}_0^{\text{MLE}}) \rightarrow_p \Sigma,$$

where Σ is a positive definite matrix.

THEOREM 3.1. *Under regularity conditions (A.1)-(A.3), if $\lambda_n \rightarrow \infty$ in a fashion such that $\lambda_n = o(\sqrt{n})$ as n goes to $+\infty$, then $P(\widehat{\beta}_j^{\text{SVS}} = 0) \rightarrow 1$ for any $j \notin \mathcal{A}$, and $\widehat{\beta}_j^{\text{SVS}} - \beta_j = O_p(n^{-1/2})$ if $j \in \mathcal{A}$ provided that the true model satisfies the same heredity principles.*

4. Computation. Similar to the original nonnegative garrote, the proposed structured variable selection and estimation procedure proceeds in two steps. First the solution path indexed by the tuning parameter M is constructed. The second step, oftentimes referred to as tuning, selects the final estimate on the solution path.

4.1. *Linear Regression.* We begin with linear regression. For both types of heredity principles, the shrinkage factors for a given M can be obtained from solving a quadratic program of the following form:

$$(4.1) \quad \min_{\theta} \left(\frac{1}{2} \|Y - Z\theta\|^2 \right), \quad \text{subject to} \quad \sum_{j=1}^p \theta_j \leq M, \text{ and } H\theta \succeq \mathbf{0},$$

where H is a $m \times p$ matrix determined by the type of heredity principles, $\mathbf{0}$ is a vector of zeros, and \succeq means “greater than or equal to” in an element-wise manner. (4.1) can be solved efficiently using standard quadratic programming techniques, and the solution path of the proposed structured variable selection and estimation procedure can be approximated by solving (4.1) for a fine grid of M 's.

Recently, Yuan and Lin (2006; 2007) showed that the solution path of the original nonnegative garrote is piecewise linear, and used this to construct an efficient algorithm for building its whole solution path. The original nonnegative garrote corresponds to the situation where the matrix H of (4.1) is a $p \times p$ identity matrix. Similar results can be expected for more general scenarios including the proposed procedures, but the algorithm will become considerably more complicated and running quadratic programming for a grid of tuning parameter tends to be a computationally more efficient alternative.

Write $\hat{B}^{\text{LS}} = \text{diag}(\hat{\beta}_1^{\text{LS}}, \dots, \hat{\beta}_p^{\text{LS}})$. The objective function of (4.1) can be expressed as

$$\frac{1}{2} \|Y - Z\theta\|^2 = \frac{1}{2} \left(Y'Y - 2\theta' \hat{B}X'Y + \theta' \hat{B}X'X \hat{B}\theta \right).$$

Because $Y'Y$ does not depend on θ 's, (4.1) is equivalent to

$$(4.2) \quad \min_{\theta} \left(-\theta' \hat{B} X' Y + \frac{1}{2} \theta' \hat{B} X' X \hat{B} \theta \right), \quad \text{subject to} \quad \sum_{j=1}^p \theta_j \leq M, \text{ and } H\theta \succeq \mathbf{0},$$

which depends on the sample size n only through $X'Y$ and the Gram matrix $X'X$. Both quantities are already computed in evaluating the least squares. Therefore, when compared with the ordinary least squares estimator, the additional computational cost of the proposed estimating procedures is free of sample size n .

Once the solution path is constructed, our final estimate is chosen on the solution path according to certain criterion. Such criterion often reflects the prediction accuracy, which depends on the unknown parameters and needs to be estimated. A commonly used criterion is the multi-fold cross validation (CV). Multi-fold CV can be used to estimate the prediction error of an estimator. The data $\mathcal{L} = \{(y_i, \mathbf{x}_i) : i = 1, \dots, n\}$ are first equally split into V subsets $\mathcal{L}_1, \dots, \mathcal{L}_V$. Using the proposed method, and data $L^{(v)} = \mathcal{L} - \mathcal{L}_v$, construct estimate $\hat{\beta}^{(v)}(M)$. The CV estimate of the prediction error is

$$\widehat{PE}(\hat{\beta}(M)) = \sum_v \sum_{(y_i, \mathbf{x}_i) \in \mathcal{L}_v} \left(y_i - \mathbf{x}_i \hat{\beta}^{(v)}(M) \right)^2$$

We select the tuning parameter M by minimizing $\widehat{PE}(\hat{\beta}(M))$. It is often suggested to use $V = 10$ in practice (Breiman, 1995).

It is not hard to see that $\widehat{PE}(\hat{\beta}^{(v)}(M))$ estimates

$$PE(\hat{\beta}(M)) = n\sigma^2 + n \left(\beta - \hat{\beta}(M) \right)' E(X'X) \left(\beta - \hat{\beta}(M) \right),$$

Since the first term is the inherent prediction error due to the noise, one often measures the goodness of an estimator using only the second term, referred to as the model error:

$$(4.3) \quad ME(\widehat{\beta}(M)) = \left(\beta - \widehat{\beta}(M)\right)' E(X'X) \left(\beta - \widehat{\beta}(M)\right),$$

Clearly, we can estimate the model error as $\widehat{PE}(\widehat{\beta}(M))/n - \widehat{\sigma}^2$ where $\widehat{\sigma}^2$ is the noise variance estimate obtained from the ordinary least squares estimate using all predictors.

4.2. *Generalized Regression.* Similarly for more general regression settings, we solve

$$(4.4) \quad \min_{\theta_0, \theta} L(Y, Z\theta + \theta_0), \quad \text{subject to} \quad \sum_{j=1}^p \theta_j \leq M, \text{ and } H\theta \succeq \mathbf{0},$$

for some matrix H . This can be done in an iterative fashion provided that the loss function L is strictly convex in its second argument. At each iteration, denote $(\theta_0^{[0]}, \theta^{[0]})$ the estimate from the previous iteration. We now approximate the objective function using a quadratic function around $(\theta_0^{[0]}, \theta^{[0]})$ and update the estimate by minimizing

$$\begin{aligned} & \sum_{i=1}^n \left(\ell'(y_i, \mathbf{z}_i \theta^{[0]} + \theta_0^{[0]}) \left[\mathbf{z}_i(\theta - \theta^{[0]}) + (\theta_0 - \theta_0^{[0]}) \right] \right. \\ & \left. + \frac{1}{2} \ell''(y_i, \mathbf{z}_i \theta^{[0]} + \theta_0^{[0]}) \left[\mathbf{z}_i(\theta - \theta^{[0]}) + (\theta_0 - \theta_0^{[0]}) \right]^2 \right) \end{aligned}$$

subject to $\sum \theta_j \leq M$, and $H\theta \succeq \mathbf{0}$, where the derivatives are taken with respect to the second argument of ℓ . Now it becomes a quadratic program.

We repeat this until certain convergence criterion is met.

In choosing the optimal tuning parameter M for general regression, we again use the multi-fold cross-validation. It proceeds in the same fashion as before except that we use a loss-dependent cross validation score:

$$\sum_v \sum_{(y_i, \mathbf{x}_i) \in \mathcal{L}_v} \ell(y_i, \mathbf{x}_i, \hat{\beta}^{(v)}(M) + \tilde{\beta}_0^{(v)}(M)).$$

5. Simulations. In this section, we investigate the finite sample properties of the proposed estimators. To fix ideas, we focus our attention on the usual normal linear regression.

5.1. *Effect of Structural Constraints.* We first consider a couple of models that represent different scenarios that may affect the performance of the proposed methods. In each of the following models, we consider three explanatory variables X_1, X_2, X_3 that follow a multivariate normal distribution with $\text{cov}(X_i, X_j) = \rho^{|i-j|}$ with three different values for ρ : 0.5, 0 and -0.5 . A quadratic model with nine terms

$$Y = \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_{11} X_1^2 + \beta_{12} X_1 X_2 + \dots + \beta_{33} X_3^2 + \epsilon,$$

is considered. Therefore, we have a total of nine predictors including three main effects, three quadratic terms and three two-way interactions. To demonstrate the importance of accounting for potential hierarchical structure among the predictor variables, we apply the nonnegative garrote estimator that recognizes strong heredity, weak heredity and without heredity constraints. In particular, we enforce the strong heredity principle by imposing the following

constraints

$$\begin{aligned} \theta_{11} \leq \theta_1, & \quad \theta_{12} \leq \theta_1, & \quad \theta_{12} \leq \theta_2, & \quad \theta_{13} \leq \theta_1, & \quad \theta_{13} \leq \theta_3, \\ & \quad \theta_{22} \leq \theta_2, & \quad \theta_{23} \leq \theta_2, & \quad \theta_{23} \leq \theta_3, & \quad \theta_{33} \leq \theta_3. \end{aligned}$$

To enforce the weak heredity, we require that

$$\begin{aligned} \theta_{11} \leq \theta_1, & \quad \theta_{12} \leq \theta_1 + \theta_2, & \quad \theta_{13} \leq \theta_1 + \theta_3, \\ & \quad \theta_{22} \leq \theta_2, & \quad \theta_{23} \leq \theta_2 + \theta_3, & \quad \theta_{33} \leq \theta_3. \end{aligned}$$

We consider two data-generating models, one follows the strong heredity principles and the other follows the weak heredity principles:

Model I. The first model follows the strong heredity principle:

$$(5.1) \quad Y = 3X_1 + 2X_2 + 1.5X_1X_2 + \epsilon;$$

Model II. The second model is similar to Model I except that the true data generating mechanism now follows the weak heredity principle:

$$(5.2) \quad Y = 3X_1 + 2X_1^2 + 1.5X_1X_2 + \epsilon;$$

For both models, the regression noise $\epsilon \sim N(0, 3^2)$.

For each model, 50 independent observations of (X_1, X_2, X_3, Y) are collected, and a quadratic model with nine terms are analyzed. We choose the tuning parameter by ten-fold cross validation as described in the last section. Following Breiman (1995), we use the model error (4.3) as the gold standard in comparing different methods. We repeat the experiment for 1000 times for

each model and the results are summarized in Table 1. The numbers in the parentheses are the standard errors. We can see that the model errors are smaller for both weak and strong heredity models compared to the model that does not incorporate any of the heredity principles. Paired t-tests confirmed that most of the observed reductions in model error are significant at 5% level.

	No Heredity	Weak Heredity	Strong Heredity
Model I			
$\rho = 0.5$	1.79 (0.05)	1.70 (0.05)	1.59 (0.04)
$\rho = 0$	1.57 (0.04)	1.56 (0.04)	1.43 (0.04)
$\rho = -0.5$	1.78 (0.05)	1.69 (0.04)	1.54 (0.04)
Model II			
$\rho = 0.5$	1.77 (0.05)	1.61 (0.05)	1.72 (0.04)
$\rho = 0$	1.79 (0.05)	1.53 (0.04)	1.70 (0.04)
$\rho = -0.5$	1.79 (0.04)	1.68 (0.04)	1.76 (0.04)

TABLE 1

Model Error Comparisons. Model I satisfies strong heredity and Model II satisfies weak heredity.

For Model I, the nonnegative garrote that respects the strong heredity principles enjoys the best performance, followed by the one with weak heredity principles. This example demonstrates the benefit of recognizing the effect heredity. Note that the model considered here also follow the weak heredity principle, which explains why the nonnegative garrote estimator with weak heredity outperforms the one that does not enforce any heredity constraints. For Model II, the nonnegative garrote with weak heredity per-

forms the best. Interestingly, the nonnegative garrote with strong heredity performs better than the original nonnegative garrote. One possible explanation is that the reduced feasible region with strong heredity, although introduces bias, at the same time makes tuning easier.

To gain further insight, we look into the model selection ability of the structured variable selection. To separate the strength of a method and effect of tuning, for each of the simulated data, we check whether or not there is any tuning parameter such that the corresponding estimate conforms with the true model. The frequency for each method to select the right model is given in Table 2, which clearly shows that the proposed structured variable selection methods pick the right models more often than the original nonnegative garrote. Note that the strong heredity version of the method can never pick Model II correctly as it violates the strong heredity principle. We also want to point out that such comparison, although useful, needs to be understood with caution. In practice, no model is perfect and selecting an additional main effect X_2 so that Model II can satisfy strong heredity may be a much more preferable alternative to many.

	No Heredity	Weak Heredity	Strong Heredity
Model I			
$\rho = 0.5$	65.5%	71.5%	82.0%
$\rho = 0$	85.0%	86.5%	90.5%
$\rho = -0.5$	66.5%	73.5%	81.5%
Model II			
$\rho = 0.5$	65.5%	75.5%	0.00%
$\rho = 0$	83.0%	90.0%	0.00%
$\rho = -0.5$	56.5%	72.5%	0.00%

TABLE 2
Frequency of Selecting the Right Model

We also checked how effective is the ten-fold cross validation in picking the right model when it does not follow any of the heredity principles. We generated the data from the model

$$Y = 3X_1 + 2X_1^2 + 1.5X_2^2 + \epsilon,$$

where the set up for simulation remains the same as before. Note that this model does not follow any of the heredity principles. For each run, we ran the nonnegative garrote with weak heredity, strong heredity, and no heredity. We chose the best among these three estimators using ten-fold cross validation. Note that the three estimators may take different values of the tuning parameter. Among 1000 runs, 64.1% of the time, nonnegative garrote with no heredity principle was elected. In contrast, for either Model I or Model II with a similar setup, less than 10% of the time nonnegative garrote with no heredity principle was elected. This is quite a satisfactory performance.

5.2. *Effect of the Size of the Interactions.* The next example is designed to illustrate the effect of the magnitude of the interaction on the proposed methods. We use a similar setup as before but now with four main effects X_1, X_2, X_3, X_4 , four quadratic terms and six two-way interactions. The true data generating mechanism is given by

$$(5.3) \quad Y = 3X_1 + 2X_2 + 1.5X_3 + \alpha(X_1X_2 - X_1X_3) + \epsilon,$$

where $\alpha = 1, 2, 3, 4$ and $\epsilon \sim N(0, \sigma^2)$ with σ^2 chosen so that the signal-to-noise ratio is always 3:1. Similar to before, the sample size $n = 50$. Figure

3 shows the mean model error estimated over 1000 runs. We can see that the strong and weak heredity models perform better than the no heredity model and the improvement becomes more significant as the strength of the interaction effect increases.

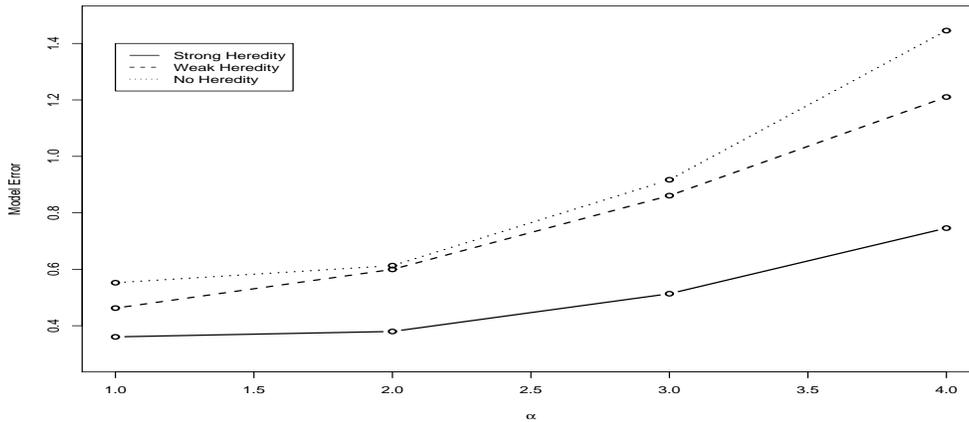


FIG 3. *Effect of the Magnitude of the Interactions*

5.3. *Large p.* To fix idea, we have focused on using least squares estimator as our initial estimator. The least squares estimators are known to perform poorly when the number of predictors is large when compared with the sample size. In particular, it is not applicable when the number of predictors exceeds the sample size. However, as shown in Yuan and Lin (2007), other initial estimators can also be used. In particular, they suggested ridge regression as one of the alternatives to the least squares estimator. To demonstrate such an extension, we consider again the regression model (5.3) but with ten main effects X_1, \dots, X_{10} , ten quadratic terms as

well as 45 interactions. The total number of effects ($p = 65$) exceeds the number of observations ($n = 50$) and therefore, the ridge regression tuned with GCV was used as the initial estimator. Figure 4 shows the solution path of the nonnegative garrote with strong heredity, weak heredity, and without any heredity for a typical simulated data with $\alpha = 4$.

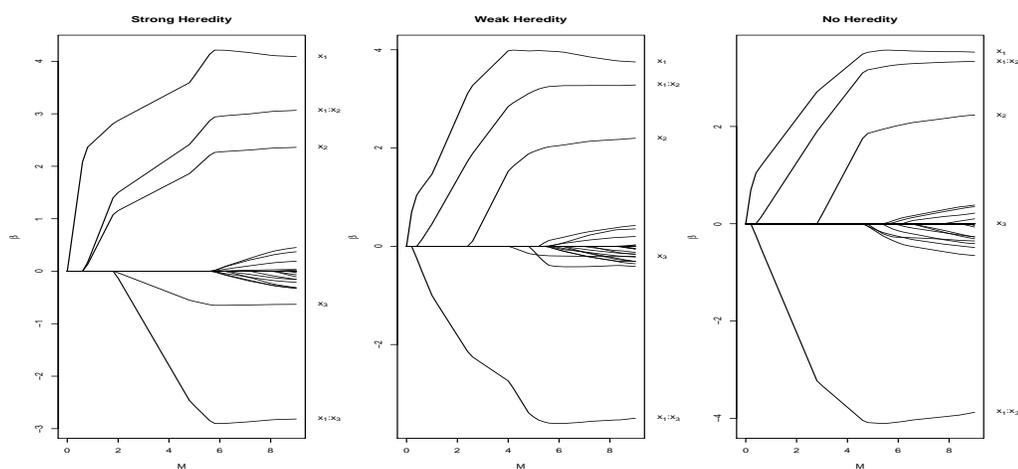


FIG 4. *Simulation when $p > n$: Solution for Different Versions of the Nonnegative Garrote*

It is interesting to notice from Figure 4 that the appropriate heredity principle, in this case strong heredity, is extremely valuable in distinguishing the true effect X_3 from other spurious effects. This further confirms the importance of heredity principles.

6. Real Data Examples. In this section, we apply the methods from Section 2 to several real data examples.

6.1. *Linear Regression Example.* The first is the prostate data, previously used in Tibshirani (1996). The data consist of the medical records of 97 male patients who were about to receive a radical prostatectomy. The response variable is the level of prostate specific antigen, and there are 8 explanatory variables. The explanatory variables are eight clinical measures: log(cancer volume) (*lcavol*), log(prostate weight) (*lweight*), age, log(benign prostatic hyperplasia amount) (*lbph*), seminal vesicle invasion (*svi*), log(capsular penetration) (*lcp*), Gleason score (*gleason*) and percentage Gleason scores 4 or 5 (*pgg45*). We consider model (1.1) with main effects, quadratic terms and two way interactions, which gives us a total of 44 predictors. Figure 5 gives the solution path of the nonnegative garrote with strong heredity, weak heredity; and without any heredity constraints. The vertical grey lines represent the models that are selected by the ten fold cross-validation.

To determine which type of heredity principle to use for the analysis, we calculated the ten-fold cross validation scores for each method. The cross-validation scores as functions of the tuning parameter M are given in the right panel of Figure 6.

Cross validation suggests the validity of heredity principles. The strong heredity is particularly favored with the smallest score. Using ten-fold cross validation, the original nonnegative garrote that neglects the effect heredity chooses a six variable model: *lcavol*, *lweight*, *lbph*, *gleason*², *lbph:svi* and *svi:pgg45*. Note that this model does not satisfy heredity principle, because

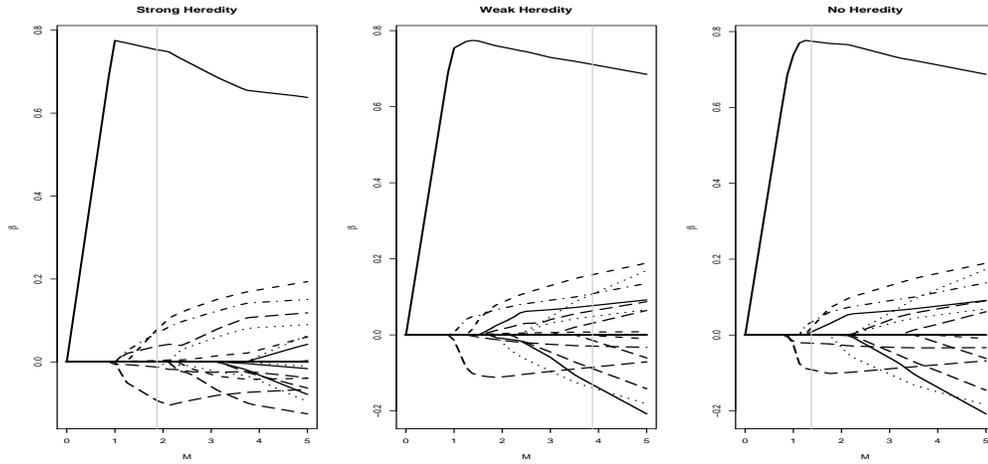


FIG 5. *Solution Path for Different Versions of the Nonnegative Garrote*

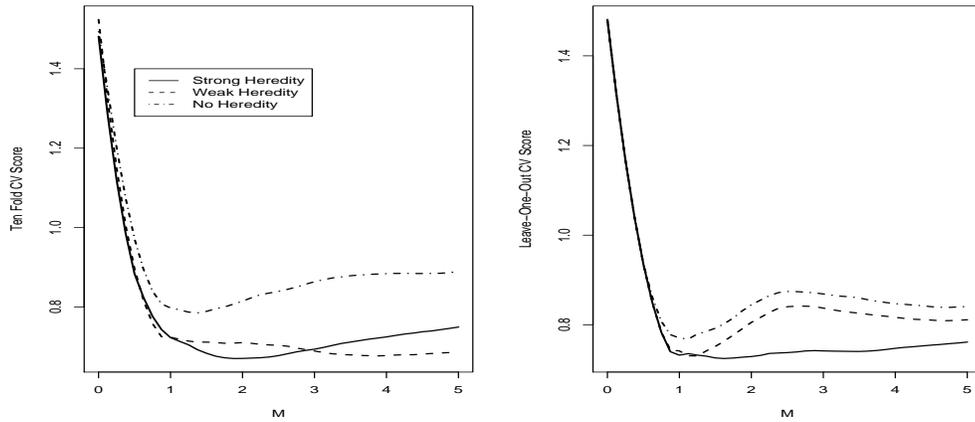
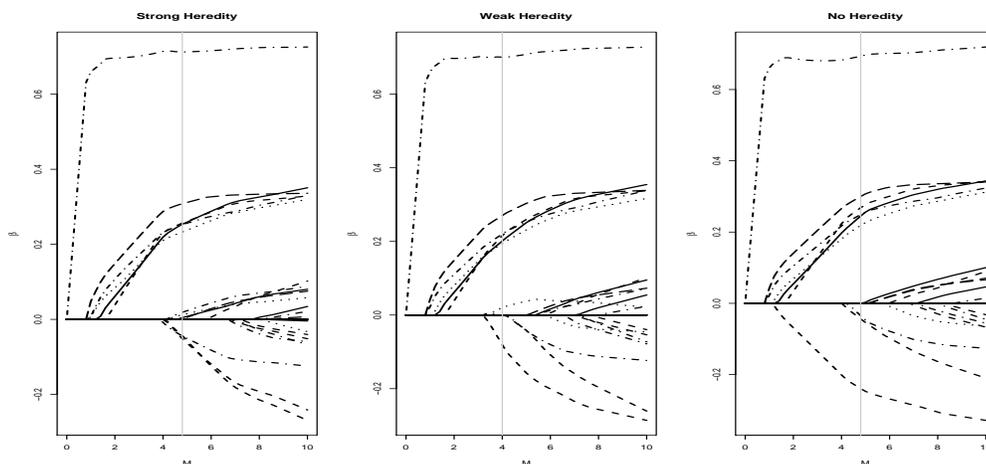


FIG 6. *Cross Validation Scores for the Prostate Data*

$gleason^2$ and $svi:pgg45$ are included without any of its parent factors. In contrast, the nonnegative garrote with strong heredity selects a model with seven variables: $lcavol$, $lweight$, $lbph$, svi , $gleason$, $gleason^2$ and $lbph:svi$. The model selected by the weak heredity, although comparable in terms of cross validation score, is considerably bigger with 16 variables. The estimated model errors for the strong heredity, weak heredity and no heredity nonnegative garrote are 0.18, 0.19 and 0.30 respectively, which clearly favors the methods that account for the effect heredity.

To further assess the variability of the ten fold cross validation, we also ran the leave-one-out cross validation on the data. The leave-one-out scores are given in the right panel of Figure 6. It shows similar pattern as the ten fold cross validation. In what follows, we shall continue to use the ten fold cross validation because of the tremendous computational advantage it brings about.

6.2. *Logistic Regression Example.* To illustrate the strategy in more general regression settings, we consider a logistic regression for the South Africa heart disease data previously used in Hastie, Tibshirani and Friedman (2003). The data consist of 9 different measures of 462 subjects and the responses indicating the presence of heart disease. We again consider a quadratic model. There is one binary predictor which leaves a total of 53 terms. Nonnegative garrote with strong heredity, weak heredity and without heredity were applied to the dataset. The solution paths are given in Figure 7.

FIG 7. *Solution Paths for the Heart Data*

The cross-validation scores for the three different methods are given in Figure 8. As we can see from the figure, nonnegative garrote with strong heredity principles achieves the lowest cross-validation score, followed by the one without heredity principles.

6.3. *Prediction Performance on Several Benchmark Data.* To gain further insight on the merits of the proposed structured variable selection and estimation techniques, we apply them to seven benchmark data sets including the previous two examples. The Ozone data, originally used in Breiman and Friedman (1985), consist of the daily maximum one-hour-average ozone reading and eight meteorological variables in the Los Angeles basin for 330 days in 1976. The goal is to predict the daily maximum one-hour-average ozone reading using the other eight variables. The Boston housing data include statistics for 506 census tracts of Boston from the 1970 census (Har-

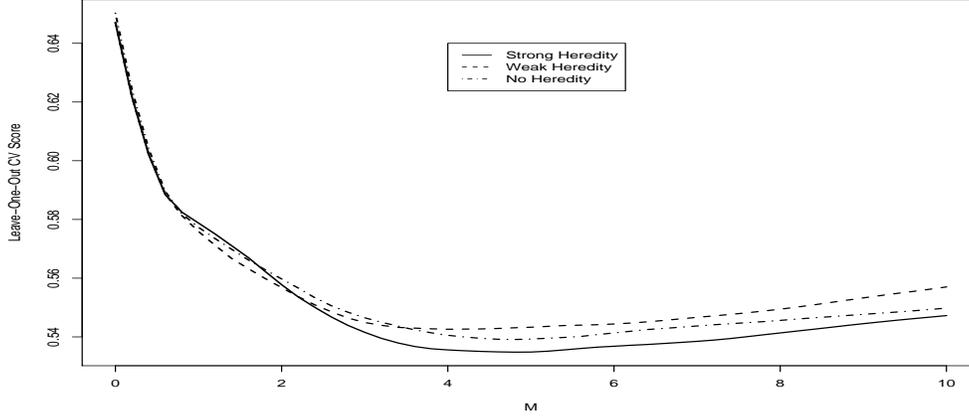


FIG 8. *Cross Validation Scores for the Heart Data*

risson and Rubinfeld, 1979). The problem is to predict the median value of owner-occupied homes based on 13 demographic and geological measures. The Diabetes data, previously analyzed by Efron et al. (2004), consist of eleven clinical measurements for a total of 442 diabetes patients. The goal is to predict a quantitative measure of disease progression one year after the baseline using the other ten measures that were collected at the baseline. Along with the prostate data, these data sets are used to demonstrate our methods in the usual normal linear regression setting.

To illustrate the performance of the structured variable selection and estimation in more general regression settings, we include two other logistic regression examples along with the South African Heart data. The Pima Indians Diabetes data have 392 observations on nine variables. The purpose is to predict whether or not a particular subject has diabetes using eight

remaining variables. The BUPA Liver Disorder data include eight variables and the goal is to relate a binary response with seven clinical measurements. Both data sets are available from the UCI Repository of machine learning databases (Newman et al., 1998).

We consider methods that do not incorporate heredity principles, respect weak or strong heredity principles. For each method, we estimate the prediction error using ten-fold cross validation, i.e., the mean squared error in the case of the four linear regression examples, and the misclassification rate in the case of the three classification examples. Table 3 documents the findings. Similar to the Heart data we discussed earlier, the total number of effects (p) can be different for the same number of main effects (q) due to the existence of binary variables. As the results from Table 3 suggest, incorporating the heredity principles leads to improved prediction for all seven data sets. Note that for the four regression datasets, the prediction error depends on the scale of the response and therefore should not be compared across datasets. For example, the response variable of the diabetes data ranges from 25 to 346 with a variance of 5943.331. In contrast, the response variable of the prostate data ranges from -0.43 to 5.48 with a variance of 1.46.

7. Discussions. When a large number of variables are entertained, variable selection becomes important. With a number of competing models that are virtually indistinguishable in fitting the data, it is often advocated to select a model with the smaller number of variables. But this principle alone may lead to models that are not interpretable. In this paper, we proposed

Data	n	q	p	No Heredity	Weak Heredity	Strong Heredity
Boston	506	13	103	12.609	12.403	12.661
Diabetes	442	10	64	3077.471	2987.447	3116.989
Ozone	203	9	54	16.558	15.100	15.397
Prostate	97	8	44	0.624	0.632	0.584
BUPA	345	6	27	0.287	0.279	0.267
Heart	462	9	53	0.286	0.275	0.262
Pima	392	8	44	0.199	0.214	0.196

TABLE 3
Prediction Performance on Seven Real Data Sets

structured variable selection and estimation methods that can effectively incorporate the hierarchical structure among the predictors in variable selection and regression coefficient estimation. The proposed methods select models that satisfy the heredity principle and are much more interpretable in practice. The proposed methods adopt the idea of the nonnegative garrote and inherit its advantages. They are easy to compute and enjoy good theoretical properties.

Similar to the original nonnegative garrote, the proposed method involves the choice of a tuning parameter which also amounts to the selection of a final model. Throughout the paper, we have focused on using the cross-validation for such purpose. Other tuning methods could also be used. In particular, it is known that prediction-based tuning may result in unnecessarily large models. Several heuristic methods are often adopted in practice to alleviate such problems. One of the most popular choice is the so-called one standard error rule (Breiman, Friedman, Stone, and Olshen 1984), where instead of choosing the model that minimizes the cross-validation score, one chooses the simplest model with a cross-validation score within one standard error

from the smallest. Our experience also suggests that a visual examination of the solution path and the cross-validation scores often leads to further insights.

The proposed method can also be used in other statistical problems whenever the structures among predictors should be respected in model building. In some applications, certain predictor variables may be known apriori to be more important than the others. This may happen, for example, in time series prediction where more recent observations generally should be more predictive of future observations.

References.

- Boyd, S. and Vandenberghe, L. (2004), *Convex Optimization*, Cambridge: Cambridge University Press.
- Breiman, L. (1995), Better subset regression using the nonnegative garrote, *Technometrics*, **37**, 373-384.
- Breiman L. and Friedman, J. (1985), Estimating optimal transformations for multiple regression and correlation, *Journal of the American Statistical Association*, **80**, 580-598.
- Breiman, L., Friedman, J., Stone, C., and Olshen, R. (1984), *Classification and Regression Trees*, New York: Chapman & Hall/CRC.
- Chipman, H. (1996), Bayesian variable selection with related predictors, *Canadian Journal of Statistics*, **24**, 17-36.
- Chipman, H., Hamada, M. and Wu, C.F.J. (1997) A Bayesian variable selection approach for analyzing designed experiments with complex aliasing, *Technometrics*, **39**, 372-381.
- Choi, N, Li, W. and Zhu, J. (2006) Variable selection with the strong heredity constraint and its oracle property, Technical report.
- Efron, B., Johnstone, I., Hastie, T. and Tibshirani, R. (2004), Least angle regression (with discussions), *The Annals of Statistics*, **32**, 407-499.
- Fan, J. and Li, R. (2001), Variable selection via nonconcave penalized likelihood and its oracle properties, *Journal of the American Statistics Association*, **96**, 1348-1360.

- George, E. I. and McCulloch, R. E. (1993), Variable selection via Gibbs sampling, *Journal of the American Statistics Association*, **88**, 881-889.
- Hamada, M. and Wu, C.F.J. (1992), Analysis of designed experiments with complex aliasing, *Journal of Quality Technology*, **24**, 130-137.
- Harrison, D. and Rubinfeld, D. (1978), Hedonic prices and the demand for clean air, *Journal of Environmental Economics and Management*, **5**, 811-820.
- Hastie, T., Tibshirani, R. and Friedman, J. (2003), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer, New York.
- Joseph, V. R. and Delaney, J. D. (2007), Functionally Induced Priors for the Analysis of Experiments. *Technometrics*, **49**, 1-11.
- Li, X., Sundarsanam, N. and Frey, D. (2006), Regularities in data from factorial experiments, *Complexity*, **11**, 32-45.
- McCullagh, P. (2002), What is a statistical model (with discussions), *Annals of Statistics*, **30**, 1225-1310.
- McCullagh, P. and Nelder, J. (1989), *Generalized Linear Models*, 2nd ed. Chapman and Hall, London.
- Nelder, J. (1977), A reformulation of linear models, *Journal of the Royal Statistics Society Series A*, **140**, 48-77
- Nelder, J. (1994), The statistics of linear models, *Statistics and Computing*, **4**, 221-234.
- Nelder, J. (1998), The selection of terms in response-surface models – how strong is the weak-heredity principle? *The American Statistician*, **52**, 315-318.
- Newman, D., Hettich, S., Blake, C. and Merz, C. (1998), UCI Repository of machine learning databases (<http://www.ics.uci.edu/~mllearn/MLRepository.html>), Irvine, CA: University of California, Department of Information and Computer Science.
- Osborne, M., Presnell, B. and Turlach, B. (2000), A new approach to variable selection in least squares problems, *IMA Journal of Numerical Analysis*, **20**, 389-403.
- Tibshirani, R. (1996), Regression shrinkage and selection via the lasso, *Journal of the Royal Statistics Society Series B*, **58**, 267-288.
- Turlach, B. (2004), Discussion of “Least angle regression”, *The Annals of Statistics*, **32**, 481-490.
- Yuan, M., Joseph, V. R. and Lin, Y. (2007), An efficient variable selection approach for analyzing designed experiments, *Technometrics*, **49**, 430-439.
- Yuan, M. and Lin, Y. (2006), Model selection and estimation in regression with grouped variables, *Journal of the Royal Statistics Society Series B*, **68**, 49-67.
- Yuan, M. and Lin, Y. (2007), On the the nonnegative garrote estimator, *Journal of the Royal Statistics Society Series B*, **69**, 143-161.

Zhao, P., Rocha, G. and Yu, B. (2006), Grouped and hierarchical model selection through composite absolute penalties, *Annals of Statistics*, to appear.

SCHOOL OF INDUSTRIAL AND SYSTEMS ENGINEERING	SCHOOL OF STATISTICS
GEORGIA INSTITUTE OF TECHNOLOGY	UNIVERSITY OF MINNESOTA
755 FERST DRIVE NW, ATLANTA, GA, 30332	224 CHURCH ST. SE, MINNEAPOLIS, MN 55455
E-MAIL: myuan@isye.gatech.edu	E-MAIL: hzou@stat.umn.edu
roshan@isye.gatech.edu	