

DEPARTMENT OF STATISTICS

University of Wisconsin

1300 University Ave.

Madison, WI 53706

TECHNICAL REPORT NO. 1141

January 2, 2008

## LASSO-Patternsearch Algorithm with Application to Ophthalmology and Genomic Data

Weiliang Shi<sup>1</sup>

Department of Statistics, University of Wisconsin, Madison WI

Grace Wahba<sup>1</sup>

Department of Statistics, Department of Computer Science and Department of Biostatistics and  
Medical Informatics, University of Wisconsin, Madison WI

Stephen Wright<sup>2</sup>

Department of Computer Science, University of Wisconsin, Madison WI

Kristine Lee<sup>3</sup>, Ronald Klein<sup>4</sup>, Barbara Klein<sup>4</sup>

Department of Ophthalmology and Visual Science, University of Wisconsin, Madison WI

---

<sup>1</sup>Research supported in part by NIH Grant EY09946, NSF Grants DMS-0505636, DMS-0604572 and ONR Grant N0014-06-0095

<sup>2</sup>Research supported in part by NSF Grants SCI-0330538, DMS-0427689, CCF-0430504, CTS-0456694, CNS-0540147 and DOE Grant DE-FG02-04ER25627

<sup>3</sup>Research supported in part by NIH grants EY06594 and EY015286

<sup>4</sup>Research support in part by NIH grants EY06594, EY015286 and Research to Prevent Blindness Senior Investigator Awards

# LASSO-Patternsearch Algorithm with Application to Ophthalmology and Genomic Data

Weiliang Shi, Grace Wahba, Stephen Wright,  
Kristine Lee, Ronald Klein and Barbara Klein  
University of Wisconsin - Madison

January 3, 2008

## Abstract

The LASSO-Patternsearch algorithm is proposed as a two-step method to identify clusters or patterns of multiple risk factors for outcomes of interest in demographic and genomic studies. The predictor variables are dichotomous or can be coded as dichotomous. Many diseases are suspected of having multiple interacting risk factors acting in concert, and it is of much interest to uncover higher order interactions or risk patterns when they exist. The patterns considered here are those that arise naturally from the log linear expansion of the multivariate Bernoulli density. The method is designed for the case where there is a possibly very large number of candidate patterns but it is believed that only a relatively small number are important. A LASSO is used to greatly reduce the number of candidate patterns, using a novel computational algorithm that can handle an extremely large number of unknowns simultaneously. Then the patterns surviving the LASSO are further pruned in the framework of (parametric) generalized linear models. A novel tuning procedure based on the GACV for Bernoulli outcomes, modified to act as a model selector, is used at both steps. We first applied the method to myopia data from the population-based Beaver Dam Eye Study, exposing physiologically interesting interacting risk factors. In particular, we found that for an older cohort the risk of progression of myopia for smokers is reduced by taking vitamins while the risk for non-smokers is independent of the “taking vitamins” variable. This is

in agreement with the general result that smoking reduces the absorption of vitamins, and certain vitamins have been associated with eye health. We then applied the method to data from a generative model of Rheumatoid Arthritis based on Problem 3 from the Genetic Analysis Workshop 15, successfully demonstrating its potential to efficiently recover higher order patterns from attribute vectors of length typical of genomic studies.

## 1 Introduction

We consider the problem which occurs in demographic and genomic studies when there are a large number of risk factors that potentially interact in complicated ways to induce elevated risk. The goal is to search for important patterns of multiple risk factors among a very large number of candidate patterns, with results that are easily interpretable. In this work the LASSO-Patternsearch algorithm (LPS) is proposed for this task. All variables are binary, or have been dichotomized before the analysis, at the risk of some loss of information; this allows the study of much higher order interactions than would be possible with risk factors with more than two possible values. Thus LPS may, if desired, be used as a preprocessor to select clusters of variables that are later analyzed in their pre-dichotomized form, see [36]. Along with demographic studies, a particularly promising application of LPS is to the analysis of patterns or clusters of SNPs (Single Nucleotide Polymorphisms) or other genetic variables that are associated with a particular phenotype, when the attribute vectors are very large and there exists a very large number of candidate patterns. LPS is designed specifically for the situation where the number of candidate patterns may be very large, but the solution, which may contain high order patterns, is believed to be sparse. LPS is based on the log linear parametrization of the multivariate Bernoulli distribution [33] to generate all possible patterns, if feasible, or at least a large subset of all possible patterns up to some maximum order. LPS begins with a LASSO algorithm (penalized Bernoulli likelihood with an  $l_1$  penalty), used with a new tuning score, BGACV. BGACV is a modified version of the GACV score [34] to target variable selection, as opposed to Kullback-Liebler distance, which is the GACV target. A novel numerical algorithm is developed specifically for this step, which can handle an extremely large number of basis functions (patterns) simultaneously. This is in particular contrast to most of the literature in the area, which uses greedy or sequential algorithms. The patterns surviving this process are then entered into a parametric linear logistic regression to obtain the final model, where further sparsity may be enforced via a

backward elimination process using the BGACV score as a stopping criterion. Properties of LPS will be examined via simulation, and in demographic data by scrambling responses to establish false pattern generation rates.

There are many approaches that can model data with binary covariates and binary responses, see, for example CART [1], LOTUS [2], Logic regression [25] and Stepwise Penalized Logistic Regression (SPLR) [24]. Logic regression is an adaptive regression methodology that constructs predictors as Boolean combinations of binary covariates. It uses simulated annealing to search through the high dimensional covariate space and uses five-fold cross validation and randomization based hypothesis testing to choose the best model size. SPLR is a variant of logistic regression with  $l_2$  penalty to fit interaction models. It uses a forward stepwise procedure to search through the high dimensional covariate space and the model size is chosen by AIC or BIC. For Gaussian data the LASSO was proposed in [30] as a variant of linear least squares ridge regression with many predictor variables. As proposed there, the LASSO minimized the residual sum of squares subject to a constraint that the sum of absolute values of the coefficients of the basis functions be less than some constant, say  $t$ . This is equivalent to minimizing the residual sum of squares plus a penalty which is some multiple  $\lambda$  (depending on  $t$ ) of the sum of absolute values ( $l_1$  penalty). It was demonstrated there that this approach tended to set many of the coefficients to zero, resulting in a sparse model, a property not generally obtaining with quadratic penalties. A similar idea was exploited in [3] to select a good subset of an overcomplete set of nonorthogonal wavelet basis functions. The asymptotic behavior of LASSO type estimators was studied in [15], and [23] discussed computational procedures in the Gaussian context. More recently [5] discussed variants of the LASSO and methods for computing the LASSO for a continuous range of values of  $\lambda$  in the Gaussian case. Variable selection properties of the LASSO were examined in [19] in some special cases, and many applications can be found on the web. In the context of nonparametric ANOVA decompositions [36] used an overcomplete set of basis functions obtained from a Smoothing Spline ANOVA model, and used  $l_1$  penalties on the coefficients of main effects and low order interaction terms, in the spirit of [3]. The present paper uses some ideas from [36], although the basis function set here is quite different. Other work has implemented  $l_1$  penalties along with quadratic (reproducing kernel square norm) penalties to take advantage of the properties of both kinds of penalties, see for example [12] [18] [35][37].

The rest of the article is organized as follows. In Section 2 we describe the first (LASSO) step of the LPS including choosing the smoothing parameter by the B-type Generalized

Approximate Cross Validation (BGACV), “B” standing for the prior belief that the solution is sparse, analogous to BIC. An efficient algorithm for the LASSO step is presented here. Section 3 describes the second step of the LASSO-Patternsearch algorithm, utilizing a parametric logistic regression, again tuned by BGACV. Section 4 presents three simulation examples, designed to demonstrate the properties of LPS as well as comparing LPS to Logic regression and SPLR. Favorable properties of LPS are exhibited in models with high order patterns and correlated attributes. Section 5 applies the method to progression of myopia in an older cohort from the Beaver Dam Eye Study [14], where some interesting risk patterns including one involving smoking and vitamins are found. Section 6 applies the method to data from a generative model of Rheumatoid Arthritis Single Nucleotide Polymorphisms adapted from the Genetics Analysis Workshop 15 [6], which examines the ability of the algorithm to recover third order patterns from extremely large attribute vectors. Section 7 notes some generalizations, and, finally, Section 8 gives a summary and conclusions. Appendix A relates the patterns to the log linear expansion of the multivariate Bernoulli distribution; Appendix B derives the BGACV score; Appendix C gives details of the specially designed code for the LASSO which is capable of handling a very large number of patterns simultaneously; Appendix D shows the detailed results of Simulation Example 3. When all of the variables are coded as 1 in the risky direction, the model will be sparsest among equivalent models. Appendix E gives a lemma describing what happens when some of the variables are coded with the opposite direction as 1.

## 2 The LASSO-Patternsearch Algorithm

### 2.1 The LASSO-Patternsearch Algorithm - Step 1

Considering  $n$  subjects, for which  $p$  variables are observed, we first reduce continuous variables to “high” or “low” in order to be able to examine *very many* variables and *their interactions* simultaneously. We will assume that for all or most of the the  $p$  variables, we know in which direction they are likely to affect the outcome or outcomes of interest, if at all. For some variables, for example smoking, it is clear for most endpoints in which direction the smoking variable is likely to be “bad” if it has any effect, and this is true of many but not all variables. For some continuous variables, for example systolic blood pressure, higher is generally “worse”, but extremely low can also be “bad”. For continuous variables, we need

to initially assume the location of a cut point on one side of which the variable is believed to be “risky” (“high”) and the other side “not risky” (“low”). For systolic blood pressure that might, for example, be 140 mmHg. For an economic variable, that might be something related to the current standard for poverty level. If the “risky” direction is known for most variables the results will be readily interpretable. Each subject thus has an attribute vector of  $p$  zeroes and ones, describing whether each of their  $p$  attributes is on one side or the other of the cutoff point. The LASSO-Patternsearch approach described below is able to deal with high order interactions and very large  $p$ . The data is  $\{y_i, x(i), i = 1, \dots, n\}$ , where  $y_i \in \{0, 1\}$  codes the response,  $x(i) = (x_1(i), x_2(i), \dots, x_p(i))$  is the attribute vector for the  $i$ th subject,  $x_j(i) \in \{0, 1\}$ . Define the basis functions  $B_{j_1 j_2 \dots j_r}(x) = \prod_{\ell=1}^r x_{j_\ell}$ , that is,  $B_{j_1 j_2 \dots j_r}(x) = 1$  if  $x_{j_1}, \dots, x_{j_r}$  are all 1’s and 0 otherwise. We will call  $B_{j_1 j_2 \dots j_r}(x)$  an  $r$ th order pattern. Let  $q$  be the highest order we consider. Then there will be  $N_B = \sum_{\nu=0}^q \binom{p}{\nu}$  patterns. If  $q = p$ , we have a complete set of  $N_B = 2^p$  such patterns (including the constant function  $\mu$ ), spanning all possible patterns. See Appendix A for a discussion of how these patterns arise from the log linear parametrization of the multivariate Bernoulli distribution. If  $q = 1$  only first order patterns (henceforth called “main effects”) are considered, if  $q = 2$  main effects and second order patterns are considered, and so forth. Letting  $p(x) = \text{Prob}[y = 1|x]$  and the logit (log odds ratio) be  $f(x) = \log[p(x)/(1 - p(x))]$ , we estimate  $f$  by minimizing

$$I_\lambda(y, f) = \mathcal{L}(y, f) + \lambda J(f) \quad (1)$$

where  $\mathcal{L}(y, f)$  is  $\frac{1}{n}$  times the negative log likelihood:

$$\mathcal{L}(y, f) = \frac{1}{n} \sum_{i=1}^n [-y_i f(x(i)) + \log(1 + e^{f(x(i))})] \quad (2)$$

with

$$f(x) = \mu + \sum_{\ell=1}^{N_B-1} c_\ell B_\ell(x), \quad (3)$$

where we are relabeling the  $N_B - 1$  (non-constant) patterns from 1 to  $N_B - 1$ , and

$$J(f) = \sum_{\ell=1}^{N_B-1} |c_\ell|. \quad (4)$$

In Step 1 of the LASSO-Patternsearch we minimize (1) using the BGACV score to choose  $\lambda$ . The next section describes the BGACV score and the kinds of results it can be expected to produce.

## 2.2 B-type Generalized Approximate Cross Validation (BGACV)

The tuning parameter  $\lambda$  in (1) balances the trade-off between data fitting and the sparsity of the model. The bigger  $\lambda$  is, the sparser the model. The choice of  $\lambda$  is generally a crucial part of penalized likelihood methods and machine learning techniques like the Support Vector Machine. For smoothing spline models with Gaussian data, [32] proposed ordinary leave-out-one cross validation (OCV). Generalized Cross Validation (GCV), derived from OCV, was proposed in [4][11], and theoretical properties were obtained in [20] and elsewhere. For smoothing spline models with Bernoulli data and quadratic penalty functionals, [34] derived the Generalized Approximate Cross Validation (GACV) from an OCV estimate following the method used to obtain GCV. In [36] GACV was extended to the case of Bernoulli data with continuous covariates and  $l_1$  penalties.

The derivation of the GACV begins with a leaving-out-one likelihood to minimize an estimate of the comparative Kullback-Leibler distance (CKL) between the true and estimated model distributions. The ordinary leave-out-one cross validation score for CKL is

$$CV(\lambda) = \frac{1}{n} \sum_{i=1}^n [-y_i f_{\lambda}^{[-i]}(x(i)) + \log(1 + e^{f_{\lambda}(x(i))})], \quad (5)$$

where  $f_{\lambda}$  is the minimizer of the objective function (1), and  $f_{\lambda}^{[-i]}$  is the minimizer of (1) with the  $i$ th data point left out. Through a series of approximations and an averaging step as described in Appendix B, we obtain the GACV score appropriate to the present context. It is a simple to compute special case of the GACV score in [36]:

$$GACV(\lambda) = \frac{1}{n} \sum_{i=1}^n [-y_i f_{\lambda_i} + \log(1 + e^{f_{\lambda_i}})] + \frac{1}{n} \text{tr} H \frac{\sum_{i=1}^n y_i (y_i - p_{\lambda_i})}{(n - N_{B_0})}, \quad (6)$$

here  $H = B^*(B^{*'}WB^*)^{-1}B^{*'}$ , where  $W$  is the  $n \times n$  diagonal matrix with  $i$ th element the estimated variance at  $x(i)$  ( $p_{i\lambda}(1 - p_{i\lambda})$ ) and  $B^*$  is the  $n \times N_{B_0}$  design matrix for the  $N_{B_0}$  non-zero  $c_{\ell}$  in the model. The quantity  $\text{tr} H \frac{\sum_{i=1}^n y_i (y_i - p_{\lambda_i})}{(n - N_{B_0})}$  plays the role of degrees of freedom for signal. As is clear from the preceding discussion, the GACV is a criterion whose target is the minimization of the (comparative) Kullback-Liebler distance from the estimate to the unknown “true” model. By analogy with the Gaussian case (where the predictive mean squared error and comparative Kullback-Liebler distance coincide), it is known that optimizing for predictive mean square error and optimizing for model selection *when the true model is sparse* are not in general the same thing. This is discussed in various places, for example, see [10] which discusses the relation between AIC and BIC, AIC being a predictive criterion

and BIC, which generally results in a sparser model, being a model selection criterion, with desirable properties when the “true” model is of fixed (low) dimension as the sample size gets large. See also [13], [19] and particularly our remarks at the end of Appendix B. In the AIC to BIC transformation, if  $\gamma$  is the degrees of freedom for the model, then BIC replaces  $\gamma$  with  $(\frac{1}{2} \log n)\gamma$ . By analogy we obtain a model selection criterion, BGACV, from GACV as follows. Letting  $\gamma$  be the quantity playing the role of degrees of freedom for signal in the Bernoulli- $l_1$  penalty case,

$$\gamma = \text{tr}H \frac{\sum_{i=1}^n y_i(y_i - p_{\lambda_i})}{(n - N_{B_0})}, \quad (7)$$

$\gamma$  is replaced by  $(\frac{1}{2} \log n)\gamma$  to obtain

$$BGACV(\lambda) = \frac{1}{n} \sum_{i=1}^n [-y_i f_{\lambda_i} + \log(1 + e^{f_{\lambda_i}})] + \frac{1}{n} \frac{\log n}{2} \text{tr}H \frac{\sum_{i=1}^n y_i(y_i - p_{\lambda_i})}{(n - N_{B_0})}. \quad (8)$$

We illustrate the difference of empirical performances between GACV and BGACV on a “true” model with a small number of strong patterns. Let  $(X_1^*, X_4^*)^T$ ,  $(X_2^*, X_5^*)^T$  and  $(X_3^*, X_6^*)^T$  be independently distributed from a bivariate normal distribution with mean 0, variance 1 and covariance 0.7.  $X_i = 1$  if  $X_i^* > 0$  and 0 otherwise,  $i = 1, 2, \dots, 6$ .  $X_7$  is independent of the others and takes two values  $\{1, 0\}$ , each with a probability of 0.5.  $X = (X_1, \dots, X_7)$ . The sample size  $n = 800$ . Three patterns that consist of six variables are important, and  $X_7$  is noise. The true logit is

$$f(x) = -2 + 1.5B_1(x) + 1.5B_{23}(x) + 2B_{456}(x). \quad (9)$$

This problem is very small so we chose the maximum order  $q = p = 7$ , so 127 patterns plus a constant term are entered in the trial model. We ran the simulation 100 times and the result is shown in Table 1. Both GACV and BGACV select the three important patterns perfectly, but GACV selects more noise patterns than BGACV. Note that a total of  $2^7 - 4$  noise patterns have been considered in each run. The maximum possible number in the last column is  $100 \times (2^7 - 4) = 12,400$ . Neither GACV or BGACV is doing a bad job but we will discuss a method to further reduce the number of noise patterns in Section 3. Figure 1 shows the scores of these two criteria in the first data set. BGACV selects a bigger smoothing parameter than GACV does. These scores are not continuous at the the point where a parameter becomes zero so we see jumps in the plots.



Table 1: The results of Simulation Example 1. The second through fourth columns are the appearance frequencies of the three important patterns in the 100 runs. The last column is the total appearance frequency of all other patterns. The second and third row compare GACV and BGACV within the first step of LPS. The fourth through sixth rows, which compare the full LPS with Logic regression and SPLR will be discussed in Section 4.1.

Method	$B_1$	$B_{23}$	$B_{456}$	other
GACV	100	100	100	749
BGACV	100	100	100	568
LPS	97	96	98	34
Logic	100	94	94	64
SPLR	100	90	55	426




Figure 1: Comparison of BGACV and GACV in the first data set of Simulation Example 1. The solid dots are the minima. BGACV selects a bigger  $\lambda$  than GACV does.

## 2.3 Computation

From a mathematical point of view, this optimization problem (1) is the same as the likelihood basis pursuit (LBP) algorithm in [36], but with different basis functions. The solution can easily be computed via a general constrained nonlinear minimization code such as MATLAB's `fmincon` on a desktop, for a range of values of  $\lambda$ , provided  $n$  and  $N_B$  are not too large. However, for extremely large data sets with more than a few attributes  $p$  (and therefore a large number  $N_B$  of possible basis functions), the problem becomes much more difficult to solve computationally with general optimization software, and algorithms that exploit the structure of the problem are needed. We design an algorithm that uses gradient information for the likelihood term in (1) to find an estimate of the correct active set (that is, the set of components  $c_\ell$  that are zero at the minimizer). When there are not too many nonzero parameters, the algorithm also attempts a Newton-like enhancement to the search direction, making use of the fact that first and second partial derivatives of the function in (1) with respect to the coefficients  $c_\ell$  are easy to compute analytically once the function has been evaluated at these values of  $c_\ell$ . It is not economical to compute the full Hessian (the matrix of second partial derivatives), so the algorithm computes only the second derivatives of the log likelihood function  $\mathcal{L}$  with respect to those coefficients  $c_\ell$  that appear to be nonzero at the solution. For the problems that the LASSO-Patternsearch is designed to solve, just a small fraction of these  $N_B$  coefficients are nonzero at the solution. This approach is similar to the two-metric gradient projection approach for bound-constrained minimization, but avoids duplication of variables and allows certain other economies in the implementation.

The algorithm is particularly well suited to solving the problem (1) for a number of different values of  $\lambda$  in succession; the solution for one value of  $\lambda$  provides an excellent starting point for the minimization with a nearby value of  $\lambda$ . Further details of this approach can be found in Appendix C.

## 3 The LASSO-Patternsearch Algorithm - Step 2

In Step 2 of LASSO-Patternsearch algorithm, the  $N_{B_0}$  patterns surviving Step 1 are entered into a linear logistic regression model using `glmfit` in MATLAB, and pattern selection is then carried out by the backward elimination method. We take out one of the  $N_{B_0}$  patterns at a time, fit the model with the remaining patterns and compute the tuning score. The

pattern that gives the best tuning score to the model after being taken out is removed from the model. This process continues until there are no patterns in the model. A final model is chosen from the pattern set with the best tuning score. Note that all-subset selection is not being done, since this will introduce an overly large number of degrees of freedom into the process.

If copious data is available, then a tuning set can be used to create the tuning score, but this is frequently not the case. Inspired by the tuning method in the LASSO step, we propose the BGACV score for the parametric logistic regression. The likelihood function is smooth with respect to the parameters so the robust assumption that appears in Appendix B is not needed. Other than that, the derivation of the BGACV score for parametric logistic regression follows that in Appendix B. Let  $s$  be the current subset of patterns under consideration and  $B_s$  be the design matrix. The BGACV score for logistic regression is the same as (8) with  $H = B_s(B_s'WB_s)^{-1}B_s'$ .

$$BGACV(s) = \frac{1}{n} \sum_{i=1}^n [-y_i f_{si} + \log(1 + e^{f_{si}})] + \frac{1}{n} \frac{\log n}{2} \text{tr} H \frac{\sum_{i=1}^n y_i (y_i - p_{si})}{(n - N_{B_0})}, \quad (10)$$

where  $f_{si}$  is the estimated log odds ratio for observation  $i$  and  $p_{si}$  is the corresponding probability. The BGACV scores are computed for each model that is considered in the backward elimination procedure, and the model with the smallest BGACV score is taken as the final model.

The following is a summary of the LASSO-Patternsearch algorithm:

1. Solve (1) and choose  $\lambda$  by BGACV. Keep the patterns with nonzero coefficients.
2. Put the patterns with nonzero coefficients from Step 1 into a logistic regression model and select models by the backward elimination method with the selection criterion being BGACV.

For simulated data, the results can be compared with the simulated model. For observational data, a selective examination of the data will be used to validate the results.

## 4 Simulation Studies

In this section we study the empirical performance of the LPS through three simulated examples. The first example continues with simulated data in Section 2.2. There are three

pairs of correlated variables and one independent variable. Three patterns are related to the response. The second example has only one high order pattern. The correlation within variables in the pattern is high and the correlation between variables in the pattern and other variables varies. The last example studies the performance of our method under various correlation settings. We compare LPS with two other methods, Logic regression [25] and Stepwise Penalized Logistic Regression (SPLR) [24]. We use the R package `LogicReg` to run Logic regression and the R package `stepPlr` to run SPLR. The number of trees and number of leaves in Logic regression are selected by 5-fold cross validation. The smoothing parameter in SPLR is also selected by 5-fold cross validation.

## 4.1 Simulation Example 1

In this example we have 7 variables and the sample size is 800. The true logit is  $f(x) = -2 + 1.5B_1(x) + 1.5B_{23}(x) + 2B_{456}(x)$ . The distribution of the covariates was described in Section 2.2. We simulated 100 data sets according to this model and ran all three methods on these data sets. The results are shown in the last three rows of Table 1.

Let's compare LPS with the LASSO step (third row in Table 1) first. LPS misses all three patterns a few times. However, these numbers are still very close to 100 and more importantly, LPS significantly reduced the number of noise patterns, from over 500 to 34. Here we see why a second step is needed after the LASSO step. Now let's look at LPS compared with the other two methods. Logic regression picks the first term perfectly but it doesn't do as well as LPS on the remaining two patterns. It also selects more noise patterns than LPS. SPLR does worse, especially on the last pattern. It is not surprising because this example is designed to be difficult for SPLR, which is a sequential method. In order for  $B_{456}$  to be in the model, at least one main effect of  $X_4$ ,  $X_5$  and  $X_6$  should enter the model first, say  $X_4$ . And then a second order pattern should also enter before  $B_{456}$ . It could be  $B_{45}$  or  $B_{46}$ . However, none of these lower order patterns are in the true model. This makes it very hard for SPLR to consider  $B_{456}$ , and the fact that variables in the higher order pattern are correlated with variables in the lower order patterns makes it even harder. We also notice that SPLR selects many more noise patterns than LPS and Logic regression. Because of the way it allows high order patterns to enter the model, the smallest SPLR model has 6 terms,  $B_1$ , one of  $B_2$  and  $B_3$ ,  $B_{23}$ , one main effect and one second order pattern in  $B_{456}$ , and  $B_{456}$ . Conditioning on the appearance frequencies of the important patterns, SPLR has to select

Table 2: The results of simulation example 2. The numerators are the appearance frequencies of  $B_{1234}$  and the denominators are the appearance frequencies of all noise patterns.

$\rho$	0	0.2	0.5	0.7
LPS	98/10	100/9	97/8	98/5
Logic	82/132	73/162	72/237	74/157
SPLR	53/602	60/576	57/581	58/552

at least  $90 + 2 \times 55 = 200$  noise patterns. The difference,  $426 - 200 = 226$  is still much bigger than 34 selected by LPS.

## 4.2 Simulation Example 2

We focus our attention on a high order pattern in this example. Let  $X_1^*$  through  $X_4^*$  be generated from a normal distribution with mean 1 and variance 1. The correlation between any of these two is 0.7.  $X_i = 1$  if  $X_i^* > 0$  and 0 otherwise for  $i = 1, \dots, 4$ .  $X_{i+4} = X_i$  with probability  $\rho$  and  $X_{i+4}$  will be generated from Bernoulli(0.84) otherwise for  $i = 1, \dots, 4$ .  $\rho$  takes values 0, 0.2, 0.5 and 0.7.  $X = (X_1, \dots, X_8)$ . Note that  $P(X_1 = 1) = 0.84$  in our simulation. The sample size is 2000 and the true logit is  $f(x) = -2 + 2B_{1234}(x)$ . We consider all possible patterns, so  $q = p = 8$ . We also ran this example 100 times and the results are shown in Table 2.

LPS does a very good job and it is very robust against increasing  $\rho$ , which governs the correlation between variables in the model and other variables. We selected the high order pattern almost perfectly and kept the noise patterns below 10 in all four settings. Logic regression selects the important pattern from 70 to 80 times and noise patterns over 130 times. There is a mild trend that it does worse as the correlation goes up, but the last one is an exception. SPLR is robust against the correlation but it doesn't do very well. It selects the important pattern from 50 to 60 times and noise patterns over 500 times. From this example we can see that LPS is extremely powerful in selecting high order patterns.

### 4.3 Simulation Example 3

The previous two examples have a small number of variables so we considered patterns of all orders. To demonstrate the power of our algorithm, we add in more noise variables in this example. The setting is similar to Example 2. Let  $X_1^*$  through  $X_4^*$  be generated from a normal distribution with mean 1 and variance 1. The correlation between any of these two is  $\rho_1$  and  $\rho_1$  takes values in 0, 0.2, 0.5 and 0.7.  $X_i = 1$  if  $X_i^* > 0$  and 0 otherwise,  $i = 1, 2, 3, 4$ .  $X_{i+4} = X_i$  with probability  $\rho_2$  and  $X_{i+4}$  will be generated from Bernoulli(0.84) otherwise for  $i = 1, 2, 3, 4$ .  $\rho_2$  takes values 0, 0.2, 0.5 and 0.7 also.  $X_9$  through  $X_{20}$  are generated from Bernoulli(0.5) independently.  $X = (X_1, \dots, X_{20})$ . The sample size  $n = 2000$  and the true logit is

$$f(x) = -2 + 2B_9(x) + 2B_{67}(x) + 2B_{1234}(x).$$

Unlike the previous two examples, we consider patterns up to the order of 4 because of the large number of variables. That gives us a total of  $\binom{20}{0} + \binom{20}{1} + \dots + \binom{20}{4} = 6196$  basis functions.

Figure 2 shows the appearance frequencies of the high order pattern  $B_{1234}$ . From the left plot we see that LPS dominates the other two methods. All methods are actually very robust against  $\rho_1$ , the correlation within the high order pattern. There is a huge gap between the two blue lines, which means SPLR is very sensitive to  $\rho_2$  which governs the correlation between variables in the high order pattern and others. This is confirmed by the right plot, where the blue lines decrease sharply as  $\rho_2$  increases. We see similar but milder behavior in Logic regression. This is quite natural because the problem becomes harder as the the noise variables become more correlated with the important variables. However, LPS handles this issue quite well, at least in the current setting. We see a small decrease in LPS as  $\rho_2$  goes up but those numbers are still very close to 100. The performance of these methods on the second order pattern  $B_{67}$  is generally similar but the trend is less obvious as a lower order pattern is easier for most methods. The main effect  $B_9$  is selected almost perfectly by every method in all settings. More detailed results are presented in Table 7 in Appendix D.

## 5 The Beaver Dam Eye Study

The Beaver Dam Eye Study (BDES) is an ongoing population-based study of age-related ocular disorders including cataract, age-related macular degeneration, visual impairment and




Figure 2: Appearance frequency of the high order pattern  $B_{1234}$  in simulation example 3. In the left panel, the x-axis is  $\rho_1$ .  $\rho_2$  is 0.2 for the dashed line and 0.7 for the solid line. In the right panel, the x-axis is  $\rho_2$ .  $\rho_1$  is 0.2 for the dashed line and 0.7 for the solid line. The red triangles represent LPS, the blue diamonds represent Logic regression and the green circles represent SPLR.

refractive errors. Between 1987 and 1988, a private census identified 5924 people aged 43 through 84 years in Beaver Dam, WI. 4926 of these people participated the baseline exam (BD I) between 1988 and 1990. Five (BD II), ten (BD III) and fifteen (BD IV) year follow-up data have been collected and there have been several hundred publications on this data. A detailed description of the study can be found in [14].

Myopia, or nearsightedness, is one of the most prevalent world-wide eye conditions. Myopia occurs when the eyeball is slightly longer than usual from front to back for a given level of refractive power of the cornea and lens and people with myopia find it hard to see objects at a distance without a corrective lens. Approximately one-third of the population experience this eye problem and in some countries like Singapore, more than 70% of the population have myopia upon completing college [27]. It is believed that myopia is related to various environmental risk factors as well as genetic factors. Five and ten-year changes of refraction for the BDES population were summarized in [16][17]. We will study the five-year progression of myopia on an older cohort aged 60 through 69 years. We focus on a small age group since the change of refraction differs for different age groups.

Based on [17] and some preliminary analysis we carried out on this data, we choose seven risk factors: *sex*, *inc*, *jomyop*, *catct*, *pky*, *asa* and *vtm* (sex, income, juvenile myopia, nuclear cataract, packyear, aspirin and vitamins). Descriptions and binary cut points are presented in Table 3. For most of these variables, we know which direction is bad. For example, male gender is a risk factor for most diseases and smoking is never good. The binary cut points are somewhat subjective here. Regarding *pky*, a pack a day for 30 years, for example, is a fairly substantial smoking history. *catct* has five levels of severity and we cut it at the third level. Aspirin (*asa*) and vitamin supplements (*vtm*) are commonly taken to maintain good health so we treat not taking them as risk factors. Juvenile myopia *jomyop* is assessed from self-reported age at which the person first started wearing glasses for distance. For the purposes of this study we have defined progression of myopia as a change in refraction of more than -0.75 diopters from baseline exam to the five year followup; accordingly  $y$  is assigned 1 if this change occurred and 0 otherwise. There are 1374 participants in this age group at the baseline examination, of which 952 have measurements of refraction at the baseline and the five-year follow-up. Among the 952 people, 76 have missing values in the covariates. We assume that the missing values are missing at random for both response and covariates, although this assumption is not necessarily valid. However the examination of the missingness and possible imputation are beyond the scope of this study. Our final data



consists of 876 subjects without any missing values in the seven risk factors.

Table 3: The variables in the progression of myopia example. The fourth column shows which direction is risky.

code	variable	unit	higher risk
<i>sex</i>	sex		Male
<i>inc</i>	income	\$1000	< 30
<i>jomyop</i>	juvenile myopia	age first wore glasses for distance	yes before age 21
<i>catct</i>	nuclear cataract	severity 1-5	4-5
<i>pky</i>	packyear	pack per day×years smoked	>30
<i>asa</i>	aspirin	taking/not taking	not taking
<i>vtm</i>	vitamins	taking/not taking	not taking

As the data set is small, we consider all possible patterns ( $q = 7$ ). The first step of the LASSO-Patternsearch algorithm selected 8 patterns, given in Table 4.

Table 4: Eight patterns selected at Step 1 in the progression of myopia data

Pattern	Estimate	Pattern	Estimate
1 <i>constant</i>	-2.9020	6 <i>pky × vtm</i>	0.6727
2 <i>catct</i>	1.9405	7 <i>inc × pky × vtm</i>	0.0801
3 <i>asa</i>	0.3000	8 <i>sex × inc × jomyop × asa</i>	0.8708
4 <i>inc × pky</i>	0.2728	9 <i>sex × inc × catct × asa</i>	0.2585
5 <i>catct × asa</i>	0.3958		

Figure 3 plots the coefficients of the 8 patterns plus the constant that survived Step 1 along with 90% confidence intervals. These patterns are then subject to Step 2, backward elimination, tuned via BGACV. The final model after the backward elimination step is

$$\begin{aligned}
 f = & -2.84 + 2.42 \times \textit{catct} + 1.11 \times \textit{pky} \times \textit{vtm} \\
 & + 1.98 \times \textit{sex} \times \textit{inc} \times \textit{jomyop} \times \textit{asa} + 1.15 \times \textit{sex} \times \textit{inc} \times \textit{catct} \times \textit{asa}. \quad (11)
 \end{aligned}$$

The significance levels for the coefficients of the four patterns in this model (11) can be formally computed and are, respectively 3.3340e-21, 1.7253e-05, 1.5721e-04, and 0.0428.




Figure 3: The eight patterns that survived Step 1 of LPS. The vertical bars are 90% confidence intervals based on linear logistic regression. Red dots mark the patterns that are significant at the 90% level. The orange dots are borderline cases, the confidence intervals barely covering 0.

The pattern  $pky \times vtm$  catches our attention because the pattern effect is strong and both variables are controllable. This model tells us that the distribution of  $y$ , progression of myopia conditional on  $pky = 1$  depends on  $catct$ , as well as  $vtm$  and higher order interactions, but progression of myopia conditional on  $pky = 0$  is independent of  $vtm$ . This interesting effect can easily be seen by going back to a table of the original  $catct$ ,  $pky$  and  $vtm$  data (Table 5). The denominators in the risk column are the number of persons with the given pattern and the numerators are the number of those with  $y = 1$ . The first two rows list the heavy smokers with cataract. Heavy smokers who take vitamins have a smaller risk of having progression of myopia. The third and fourth rows list the heavy smokers without cataract. Again, taking vitamins is protective. The first four rows suggest that taking vitamins in heavy smokers is associated with a reduced risk of getting more myopic. The last four rows list all non-heavy smokers. Apparently taking vitamins does not similarly reduce the risk of becoming more myopic in this population. Actually, it is commonly known that smoking significantly decreases the serum and tissue vitamin level, especially Vitamin C and Vitamin E, for example [8]. Our data suggest a possible reduction in the progression of myopia in persons who smoke who take vitamins. However, our data are observational and subject to uncontrolled confounding. A randomized controlled clinical trial would provide the best evidence of any effect of vitamins on progression of myopia in smokers.

Table 5: The raw data for cataract, smoking and not taking vitamins.

catct	pky	no vitamins	risk
1	1	1	$17/23 = 0.7391$
1	1	0	$7/14 = 0.5000$
0	1	1	$22/137 = 0.1606$
0	1	0	$2/49 = 0.0408$
1	0	1	$18/51 = 0.3529$
1	0	0	$19/36 = 0.5278$
0	0	1	$22/363 = 0.0606$
0	0	0	$13/203 = 0.0640$

Since the model is the result of previous data mining, caution in making significance statements may be in order. To investigate the probability of the overall procedure to

generate significant false patterns, we kept the attribute data fixed, randomly scrambled the response data and applied the LPS algorithm on the scrambled data. The procedure was repeated 1000 times, and in all these runs, 1 main effect, 10 second order, 5 third order and just 1 fourth order patterns showed up. We then checked on the raw data. There are 21 people with the pattern  $sex \times inc \times jomyop \times asa$  and 9 of them have progression of myopia. The incidence rate is 0.4286, as compared to the overall rate of 0.137. Note that none of the variables in this pattern are involved with variables in the two lower order patterns  $catct$  and  $pky \times vtm$  so it can be concluded that the incidence rate is contributed only by the pattern effect. People with the other size four pattern  $sex \times inc \times catct \times asa$  have an incidence rate of 0.7727 (17 out of 22), which can be compared with the incidence rate of all people with  $catct$ , 0.4919.

We also applied Logic regression and SPLR on this data set. Logic selected both  $catct$  and  $pky \times vtm$  but missed the two high order patterns. Instead, it selected  $asa$  as a main effect. Note that  $asa$  is present in both size four patterns. SPLR selected the same patterns as Logic regression with an addition of  $pky$ , which is necessary for  $pky \times vtm$  to be included. These results agree with what we have found in the simulation studies: they are not as likely as LPS in finding higher order patterns. It is noted that in the original version of LPS [29], Step 1 was tuned by GACV rather than BGACV, and resulted in the above eight patterns in Table 4 plus four more, but the final model after Step 2 was the same.

## 6 Rheumatoid Arthritis and SNPs in GAW 15

Rheumatoid arthritis (RA) is a complex disease with a moderately strong genetic component. Generally females are at a higher risk than males. Many studies have implicated a specific region on chromosome 6 as being related to the risk of RA, recently [7], although possible regions on other chromosomes have also been implicated. The 15th Genetic Analysis Workshop (GAW 15, November 2006 [6]) focused on RA, and an extensive simulation data set of cases and controls with simulated single nucleotide polymorphisms (SNPs) was provided to participants and is now publicly available [22]. SNPs are DNA sequence variations that occur when a single nucleotide in the genome sequence is changed. Many diseases are thought to result from SNP changes at multiple sites that may interact, thus it is important to have tools that can ferret out groups of possibly interacting SNPs.

We applied LPS to some of the simulated SNP RA GAW 15 data [28]. This provided an

opportunity to apply LPS in a context with large genetic attribute vectors, with a known genetic architecture, as described in [22], and to compare the results against the description of the architecture generating the data. The description of the data generator does not explicitly describe higher order interactions, and our original analysis of some of the chromosome 6 data there did not find higher order patterns. However, since our original work (submitted in November of 2006 and embargoed until July 2007) the LASSO optimization code has been further developed to have the ability to deal with an extremely large number of unknowns. Thus we decided to use the GAW data to build a simulation study where we modified the model that appears in [28] to introduce a three way pattern, and in the process deal with some anomalous minus signs in our fitted model, also observed by others [26]. We then simulated phenotype data from the GAW genotypes and covariates, and can evaluate how well the improved LPS of this paper reproduces the model generating the data with the three way pattern. This section describes the results.

In the simulated genetic data sets [22] genome wide scans of 9187 SNPs were generated with just a few of the SNPs linked to rheumatoid arthritis, according to the described model architecture. The data simulation was set up to mimic the familial pattern of rheumatoid arthritis including a strong chromosome 6 effect. A large population of nuclear families (two parents and two offspring) was generated. This population contains close to 2 million sibling pairs. From this population, a random sample of 1500 families was selected from among families with two affected offspring and another random sample of 2000 families was selected from among families where no member was affected. A total of 100 independent (replicate) data sets were generated. We randomly picked one offspring from each family in replicate 1 as our data set. As for covariates we take the 674 SNPs in chromosome 6 that were generated as a subset of the genome wide scan data and three environmental variables: *age*, *sex* and *smoking*. We created two dummy variables for each SNP since most of them have three levels: normal, one variant allele and two variant alleles. For environmental variables female gender is treated as a risk factor, smoking is treated as a risk factor and  $\text{age} \geq 55$  is treated as a risk factor. We first describe a reanalysis of this data using the LPS algorithm of this paper. We began our analysis with a screen step. In this step, each variable is entered into a linear logistic regression model. For SNPs with three levels, both dummy variables are entered into the same model. We fit these models and keep the variables with at least one  $p$ -value less than 0.05. The screen step selected 74 variables (72 SNPs plus *sex* and *smoking*). We then ran LPS on these 74 variables with  $q = 2$ , which generates 10371 basis

functions. The final model was

$$f = -.62 + .87 \times \textit{smoking} + 1.05 \times \textit{sex} - 2.04 \times \textit{SNP6\_153\_1} \\ - 1.45 \times \textit{SNP6\_154\_1} + 2.23 \times \textit{SNP6\_162\_1} - 5.60 \times \textit{SNP6\_153\_2}, \quad (12)$$

where *SNP6\_153\_1* is SNP number 153 on chromosome 6 with 1 variant allele, *SNP6\_154\_1* is SNP number 154 on chromosome 6 with 1 variant allele, *SNP6\_162\_1* is SNP number 162 on chromosome 6 with 1 variant allele and *SNP6\_153\_2* is SNP number 153 on chromosome 6 with 2 variant alleles. When the analysis in [28] was presented at the GAW 15 workshop in November 2006 the tuning procedure presented here had not been finalized. Furthermore, the algorithm at the core of the LASSO-Patternsearch could only handle 4000 patterns, so a much more stringent screen procedure was used to send only 1584 basis functions to Step 1 and both Step 1 and Step 2 were tuned against prediction accuracy using replicate 2 as a tuning set. The model (12) obtained here is exactly the same as [28]. We were pleased to find that the reanalysis here was able to handle and successfully reject a much larger number of patterns in Step 1 and that the in-sample BGACV tuning here was just as good as having a separate tuning set. It is interesting to note that in this particular problem, tuning for prediction and for model selection apparently led to the same results, although in general this is not necessarily the case.

According to the description of the architecture generating the data the risk of RA is affected by two loci (C and D) on chromosome 6, *sex*, *smoking* and a *sex* by locus C interaction. It turns out that both *SNP6\_153* and *SNP6\_154* are very close to locus C on chromosome 6 and *SNP6\_162* is very close to locus D on chromosome 6. The LPS method picked all important variables without any false positives. The description of the data generation architecture said that there was a strong interaction between *sex* and locus C. We didn't pick up a *sex* by locus C interaction, which was very surprising. However, we were able to test the prediction accuracy of this model as well as the all genome data LPS model reported in [28] because of the other independent data sets generated according to the same mechanism. Using replicate 3 as our test set, and noting that the sample sizes for treatments and controls were the same for all replicates, we predicted a subject to be diseased if the predicted probability is bigger than 0.5 and predicted him or her to be healthy if otherwise. A variety of analysis methods for the all genome data set were presented at GAW 15, and some had reported error rates, so we were able to compare the LPS analysis as applied to the all genome data set with that of some of the other analyses of the same

data set, and it was found to be competitive with the best error rate reported.

We were curious about the apparent counter-intuitive negative coefficients for both *SNP6\_153* patterns and the *SNP6\_154\_1* pattern, which appear to say that normal alleles are risky and variant alleles are protective. Others also found an anomalous protective effect for *SNP6\_154* normal alleles [26]. We went back and looked at the raw data for *SNP6\_153* and *SNP6\_154* as a check but actually Table 4 of [28] shows that this protective effect is in the simulated data, for whatever reason, and it also shows that this effect is stronger for women than for men. We then recoded the *SNP6\_153* and *SNP6\_154* responses to reflect the actual effect of these two variables as can be seen in tables of the simulated data.

Table 6: Fitted and Simulated Models, see text for explanation.

	Variable 1	Level 1	Variable 2	Level 2	Coef	Est
Main effects	constant	-	-	-	-4.8546	-4.6002
	<i>smoking</i>	-	-	-	0.8603	0.9901
	<i>SNP6_153</i>	1	-	-	1.8911	1.5604
	<i>SNP6_162</i>	1	-	-	2.2013	1.9965
	<i>SNP6_154</i>	2	-	-	0.7700	1.0808
Second order patterns	<i>sex</i>	-	<i>SNP6_153</i>	1	0.7848	0.9984
	<i>sex</i>	-	<i>SNP6_154</i>	2	0.9330	0.9464
	<i>SNP6_153</i>	2	<i>SNP6_154</i>	2	4.5877	4.2465
	<i>SNP6_153</i>	1	<i>SNP6_553</i>	2	0.4021	0
	<i>SNP6_154</i>	2	<i>SNP6_490</i>	1	0.3888	0
Added						
Third order pattern	<i>sex</i> × <i>SNP6_108_2</i> × <i>SNP6_334_2</i>				3	2.9106

Table 6 shows the results. The new fitted model, above the double line, has four main effects and five second order patterns. The estimated coefficients are given in the column headed “Coef”. The  $sex \times SNP6_153$  and  $sex \times SNP6_154$  are there as expected. Two weak second order patterns involving *SNP6\_553* and *SNP6\_490* are fitted, but do not appear to be explained by the simulation architecture. This model resulted from fitting with  $q = 2$ . Then the LPS algorithm was run to include all third order patterns ( $q = 3$ ) of the 74 variables which passed the screen step. This generated 403,594 basis functions. No third

order patterns were found, and the fitted model was the same as in the  $q = 2$  case. To see if a third order pattern would be found if it were there, we created a generative model with the four main effects and five second order patterns of Table 1, with their coefficients from the “Coef” column, and added to it a third order pattern  $sex \times SNP6_{108_2} \times SNP6_{334_2}$  with coefficient 3. The two SNPs in the third order pattern were chosen to be well separated in chromosome 6 from the reported gene loci. LPS did indeed find the third order pattern. The estimated coefficients are found under the column headed “Est”. No noise patterns were found, and the two weak second order patterns in the model were missed. However, the potential for the LPS algorithm to find higher order patterns is clear. Further investigation of the properties of the method in genotype-phenotype scenarios is clearly warranted, taking advantage of the power of the LASSO algorithm to handle a truly large number of unknowns simultaneously. Run time was 35.2 minutes on an AMD Dual-Core 2.8 GHz machine with 64 GB memory. Using multiple runs with clever designs to guarantee that every higher order pattern considered is in at least one run, will allow the analysis of much larger SNP data sets with tolerable computer cost.

## 7 Discussion

In any problem where there are a large number of highly interacting predictor variables that are or can be reduced to dichotomous values, LPS can be profitably used. If the “risky” direction (with respect to the outcome of interest) is known for all or almost all of the variables, the results are readily interpretable. If the risky direction is coded correctly for all of the variables, the fitted model can be expected to be sparser than that for any other coding. However, if a small number of risky variables are coded in the “wrong” way, this usually can be detected. The method can be used as a preprocessor when there are very many continuous variables in contention, to reduce the number of variables for more detailed nonparametric analysis.

LPS, using the algorithm of Appendix C is efficient. On an Intel Quad-core 2.66 GHz machine with 8GB memory the LPS Steps 1 and 2 can do 50,000 basis functions in 2.33 hours. On an AMD Dual-Core 2.8 GHz machine with 64 GB memory the algorithm did LPS with 403,594 basis functions in 32 minutes. It can do 2,000,000 basis functions in 9.67 hours. On the same AMD machine, problems of the size of the myopia data (128 unknowns) can be solved in a few seconds and the GAW 15 problem data (10371 unknowns) was solved in



2.3 minutes.

A number of considerations enter into the choice of  $q$ . If the problem is small and the user is interested in high order patterns, it doesn't hurt to include all possible patterns; if the problem is about the size of Simulation Example 3,  $q = 4$  might be a good choice; For genomic data the choice of  $q$  can be limited by extremely large attribute vectors. In genomic or other data where the existence of a very small number of important higher patterns is suspected, but there are too many candidates to deal with simultaneously, it may be possible to overcome the curse of dimensionality with multiple screening levels and multiple runs. For example, considering say, third or even fourth order patterns, variables could be assigned to doable sized runs so that every candidate triple or quadruple of variables is assigned to at least one run. With our purpose built algorithm, the approach is quite amenable to various flavors of exploratory data mining. When a computing system such as Condor (<http://www.cs.wisc.edu/condor/>) is available, many runs can compute simultaneously.

Many generalizations are available. Two classes of models where the LASSO-Patternsearch approach can be expected to be useful are the multicategory end points model in [21][31], where an estimate is desired of the probability of being in class  $k$  when there are  $K$  possible outcomes; another is the multiple correlated endpoints model in [9]. In this latter model, the correlation structure of the multiple endpoints can be of interest. Another generalization allows the coefficients  $c_\ell$  to depend on other variables; however, the penalty functional must involve  $\ell_1$  penalties if it is desired to have a convex optimization problem with good sparsity properties with respect to the patterns. In studies with environmental as well as genomic data selected interactions between SNP patterns and continuous covariates can be examined [36].

## 8 Summary and Conclusions

The LASSO-Patternsearch algorithm brings together several known ideas in a novel way, using a tailored tuning and pattern selection procedure and a new purpose built computational algorithm. We have examined the properties of the LPS by analysis of observational data, and simulation studies at a scale similar to the observational data. The results are verified in the simulation studies by examination of the generated "truth", and in the observational data by selective examination of the observational data directly, and data scrambling to check false alarm rates, with excellent results. The novel computational algorithm allows

the examination of a very large number of patterns, and, hence, high order interactions. We believe the LASSO-Patternsearch will be an important addition to the toolkit of the statistical data analyst.

## Acknowledgments

Thanks to David Callan for many helpful suggestions and for Appendix E. Research supported by NIH Grants EY09946 (WS, GW) NSF Grants DMS-0505636, DMS-0604572 (WS, GW) ONR Grant N0014-06-0095 (WS,GW) NIH Grants EY06594 and EY015286 (BEK, RK, KL) Research to Prevent Blindness Senior Investigator Awards (RK, BEK) NSF Grants SCI-0330538, DMS-0427689, CCF-0430504, CTS-0456694, CNS-0540147, DOE Grant DE-FG02-04ER25627 (SW)

# Appendices

## A The Multivariate Bernoulli Distribution and Pattern Basis Functions

Let  $(x_0, x_1, \dots, x_p)$  be a  $p + 1$  dimensional vector of possibly correlated Bernoulli random variables. The most general form of the joint density is

$$p(x_0, x_1, \dots, x_p) = p(0, 0, \dots, 0)^{\pi_{j=0}^p(1-x_j)} p(1, 0, \dots, 0)^{x_0 \pi_{j=1}^p(1-x_j)} \dots p(1, 1, \dots, 1)^{\pi_{j=0}^p x_j}. \quad (13)$$

Taking logs gives

$$\begin{aligned} \log p(x_0, x_1, \dots, x_p) &= \pi_{j=0}^p(1-x_j) \log p(0, 0, \dots, 0) \\ &+ x_0 \pi_{j=1}^p(1-x_j) \log p(1, 0, \dots, 0) \\ &+ \dots \\ &+ \pi_{j=0}^p x_j \log p(1, 1, \dots, 1) \end{aligned}$$

and, following [33] collecting terms gives

$$\begin{aligned}
\log p(x_0, x_1, \dots, x_p) &= \phi + \sum_{j=0}^p \mu_j x_j + \sum_{0 \leq j < k} \mu_{jk} x_j x_k \\
&+ \sum_{0 \leq j < k < \ell} \mu_{jkl} x_j x_k x_\ell \\
&+ \dots \\
&+ \mu_{0,1,\dots,p} x_0 x_1 x_2 \dots x_p.
\end{aligned}$$

The  $\mu$  terms involve various cross product ratios that can be interpreted in various contexts, and the non-zero  $\mu$ 's define a graphical model, see [33]. It is now easy to see what the log odds ratio (logit)  $f(x_1, \dots, x_p)$  for variable  $x_0$  given  $(x_1, \dots, x_p)$  is:

$$f(x_1, \dots, x_p) = \log p(1, x_1, \dots, x_p) - \log p(0, x_1, \dots, x_p), \quad (14)$$

and is easily seen to be

$$f(x_1, \dots, x_p) = \mu_0 + \sum_{j=1}^p \mu_{0j} x_j + \sum_{1 < j < k} \mu_{0jk} x_j x_k + \dots + \mu_{0,1,\dots,p} x_1 \dots x_p. \quad (15)$$

which is exactly of the form of Eqn (3) with  $N_B = 2^p - 1$ .

## B The BGACV Score

We denote the estimated logit function by  $f_\lambda(\cdot)$  and define  $f_{\lambda i} = f_\lambda(x(i))$  and  $p_{\lambda i} = \frac{e^{f_{\lambda i}}}{1 + e^{f_{\lambda i}}}$  for  $i = 1, \dots, n$ . Now define

$$OBS(\lambda) = \frac{1}{n} \sum_{i=1}^n [-y_i f_{\lambda i} + \log(1 + e^{f_{\lambda i}})]. \quad (16)$$

From (2.1) in [34], the leave-one-out CV is

$$\begin{aligned}
CV(\lambda) &= \frac{1}{n} \sum_{i=1}^n [-y_i f_{\lambda i}^{[-i]} + \log(1 + e^{f_{\lambda i}})] \\
&= OBS(\lambda) + \frac{1}{n} \sum_{i=1}^n y_i (f_{\lambda i} - f_{\lambda i}^{[-i]}) \\
&= OBS(\lambda) + \frac{1}{n} \sum_{i=1}^n y_i \frac{f_{\lambda i} - f_{\lambda i}^{[-i]}}{y_i - p_{\lambda i}^{[-i]}} \times \frac{y_i - p_{\lambda i}}{1 - (p_{\lambda i} - p_{\lambda i}^{[-i]}) / (y_i - p_{\lambda i}^{[-i]})}. \quad (17)
\end{aligned}$$

Following [34], [36] approximate  $(p_{\lambda_i} - p_{\lambda_i}^{[-i]})/(y_i - p_{\lambda_i}^{[-i]})$  by Taylor expansion.

$$\begin{aligned} \frac{p_{\lambda_i} - p_{\lambda_i}^{[-i]}}{y_i - p_{\lambda_i}^{[-i]}} &\approx \dot{p}(f_{\lambda_i}) \frac{f_{\lambda_i} - f_{\lambda_i}^{[-i]}}{y_i - p_{\lambda_i}^{[-i]}} \\ &= \frac{e^{f_{\lambda_i}}}{(1 + e^{f_{\lambda_i}})^2} \frac{f_{\lambda_i} - f_{\lambda_i}^{[-i]}}{y_i - p_{\lambda_i}^{[-i]}}. \end{aligned} \quad (18)$$

Denote

$$G_i \equiv \frac{f_{\lambda_i} - f_{\lambda_i}^{[-i]}}{y_i - p_{\lambda_i}^{[-i]}}; \quad (19)$$

then from (17) and (18), we get

$$CV(\lambda) \approx OBS(\lambda) + \frac{1}{n} \sum_{i=1}^n G_i \frac{y_i(y_i - p_{\lambda_i})}{1 - \frac{e^{f_{\lambda_i}}}{(1+e^{f_{\lambda_i}})^2} G_i}. \quad (20)$$

Denote the objective function in (1) - (4) by  $I_\lambda(y, c)$ , let  $B_{ij} = B_j(x(i))$  be the entries of the design matrix  $B$ , and, for ease of notation denote  $\mu = c_{N_B}$ . Then the objective function can be written

$$I_\lambda(y, c) = \frac{1}{n} \sum_{i=1}^n [-y_i \sum_{j=1}^{N_B} c_j B_{ij} + \log(1 + e^{(\sum_{j=1}^{N_B} c_j B_{ij})})] + \lambda \sum_{j=1}^{N_B-1} |c_j|. \quad (21)$$

Denote the minimizer of (21) by  $\hat{c}$ . We know that the  $l_1$  penalty produces sparse solutions. Without loss of generality, we assume that the first  $s$  components of  $\hat{c}$  are nonzero. When there is a small perturbation  $\epsilon$  on the response, we denote the minimizer of  $I_\lambda(c, y + \epsilon)$  by  $\tilde{c}$ . The 0's in the solutions are robust against a small perturbation in the response. That is, when  $\epsilon$  is small enough, the 0 elements will stay at 0. This can be seen by looking at the the KKT conditions when minimizing (21). Therefore, the first  $s$  components of  $\tilde{c}$  are nonzero and the rest are zero. For simplicity, we denote the first  $s$  components of  $c$  by  $c^*$  and the first  $s$  columns of the design matrix  $B$  by  $B^*$ . Then let  $f_\lambda^y$  be the column vector with  $i$  entry  $f_\lambda(x(i))$  based on data  $y$ , and let  $f_\lambda^{y+\epsilon}$  be the same column vector based on data  $y + \epsilon$ .

$$f_\lambda^{y+\epsilon} - f_\lambda^y = B(\tilde{c} - \hat{c}) = B^*(\tilde{c}^* - \hat{c}^*). \quad (22)$$

Now we take the first-order Taylor expansion of  $\frac{\partial I_\lambda}{\partial c^*}$ :

$$\left[ \frac{\partial I_\lambda}{\partial c^*} \right]_{(\tilde{c}, y+\epsilon)} \approx \left[ \frac{\partial I_\lambda}{\partial c^*} \right]_{(\hat{c}, y)} + \left[ \frac{\partial^2 I_\lambda}{\partial c^* \partial c^{*t}} \right]_{(\hat{c}, y)} (\tilde{c}^* - \hat{c}^*) + \left[ \frac{\partial^2 I_\lambda}{\partial c^* \partial y^t} \right]_{(\hat{c}, y)} (y + \epsilon - y). \quad (23)$$

Define

$$U \equiv n \left[ \frac{\partial^2 I_\lambda}{\partial c^* \partial c'^*} \right]_{(\hat{c}, y)} = B^{*'} \text{diag} \left( \left[ \frac{e^{f_{\lambda 1}}}{(1 + e^{f_{\lambda 1}})^2}, \dots, \frac{e^{f_{\lambda n}}}{(1 + e^{f_{\lambda n}})^2} \right] \right) B^* = B^{*'} W B^*, \quad \text{say,}$$

and

$$V \equiv -n \left[ \frac{\partial^2 I_\lambda}{\partial c^* \partial y'} \right]_{(\hat{c}, y)} = B^{*'}.$$

By the first-order conditions, the left-hand side and the first term of the right-hand side of (23) are zero. So we have

$$U(\tilde{c}^* - \hat{c}^*) \approx V\epsilon. \quad (24)$$

Combine (22) and (24) we have  $f_\lambda^{y+\epsilon} - f_\lambda^y \approx H\epsilon$ , where

$$H \equiv B^* U^{-1} V \equiv B^* U^{-1} B^{*'} \quad (25)$$

Now consider the following perturbation  $\epsilon_0 = (0, \dots, p_{\lambda i}^{[-i]} - y_i, \dots, 0)'$ ; then  $f_\lambda^{y+\epsilon_0} - f_\lambda^y \approx \epsilon_{0i} H_{\cdot i}$ , where  $\epsilon_{0i} = p_{\lambda i}^{[-i]} - y_i$  and  $H_{\cdot i}$  is the  $i$ th row of  $H$ . By the Leave-Out-One Lemma (stated below),  $f_\lambda^{[-i]} = f_\lambda^{y+\epsilon_0}$ . Therefore

$$G_i = \frac{f_{\lambda i} - f_{\lambda i}^{[-i]}}{y_i - p_{\lambda i}^{[-i]}} = \frac{f_{\lambda i}^{y+\epsilon_0} - f_{\lambda i}}{\epsilon_{0i}} \approx h_{ii} \quad (26)$$

where  $h_{ii}$  is the  $ii$ th entry of  $H$ . From (20), the ACV score is

$$ACV(\lambda) = \frac{1}{n} \sum_{i=1}^n [-y_i f_{\lambda i} + \log(1 + e^{f_{\lambda i}})] + \frac{1}{n} \sum_{i=1}^n h_{ii} \frac{y_i(y_i - p_{\lambda i})}{(1 - w_{ii} h_{ii})} \quad (27)$$

where we have used  $w_{ii} = \frac{e^{f_{\lambda i}}}{(1 + e^{f_{\lambda i}})^2}$ , (the estimated variance at  $x(i)$ ). The  $GACV$  score is obtained from the ACV score in (27) by replacing  $h_{ii}$  by  $\frac{1}{n} \text{tr}(H)$  and  $w_{ii} h_{ii}$  by  $\frac{1}{n} \text{tr}(WH)$ . It is not hard to see that  $\text{tr}(WH) = \text{tr} W^{1/2} H W^{1/2} = s \equiv N_{B_0}$ , the number of basis functions in the model, giving

$$GACV(\lambda) = \frac{1}{n} \sum_{i=1}^n [-y_i f_{\lambda i} + \log(1 + e^{f_{\lambda i}})] + \frac{1}{n} \text{tr} H \frac{\sum_{i=1}^n y_i(y_i - p_{\lambda i})}{(n - N_{B_0})}. \quad (28)$$

Adding the weight  $\frac{1}{2} \log n$  to the ‘‘optimism’’ part of the  $GACV$  score, we obtain the B-type  $GACV$  (BGACV):

$$BGACV(\lambda) = \frac{1}{n} \sum_{i=1}^n [-y_i f_{\lambda i} + \log(1 + e^{f_{\lambda i}})] + \frac{1}{n} \frac{\log n}{2} \text{tr} H \frac{\sum_{i=1}^n y_i(y_i - p_{\lambda i})}{(n - N_{B_0})}. \quad (29)$$

**Lemma B.1** (*Leave-Out-One Lemma*)

Let the objective function  $I_\lambda(y, f)$  be defined as before. Let  $f_\lambda^{[-i]}$  be the minimizer of  $I_\lambda(y, f)$  with the  $i$ th observation omitted and let  $p_\lambda^{[-i]}$  be the corresponding probability. For any real number  $\nu$ , we define the vector  $z = (y_1, \dots, y_{i-1}, \nu, y_{i+1}, \dots, y_n)'$ . Let  $h_\lambda(i, \nu, \cdot)$  be the minimizer of  $I_\lambda(z, f)$ ; then  $h_\lambda(i, p_\lambda^{[-i]}, \cdot) = f_\lambda^{[-i]}(\cdot)$ .

The proof of lemma B.1 is quite simple and very similar to the proof of the Leave-Out-One-Lemma in [36] so we will omit it here.

We remark that in this paper we have employed the BGACV criterion twice as a stringent model selector under the assumption that the true or the desired model is sparse. Simulation experiments (not shown) suggest that the GACV criterion is preferable if the true model is not sparse and/or the signal is weak. The GACV and the BGACV selections probably bracket the region of interest of  $\lambda$  in most applications.

## C Minimizing the Penalized Log Likelihood Function

The function (1) is not differentiable with respect to the coefficients  $\{c_\ell\}$  in the expansion (4), so most software for large-scale continuous optimization cannot be used to minimize it directly. We can however design a specialized algorithm that uses gradient information for the smooth term  $\mathcal{L}(y, f)$  to form an estimate of the correct active set (that is, the set of components  $c_\ell$  that are zero at the minimizer of (1)). Some iterations of the algorithm also attempt a Newton-like enhancement to the search direction, computed using the projection of the Hessian of  $\mathcal{L}$  onto the set of nonzero components  $c_\ell$ . This approach is similar to the two-metric gradient projection approach for bound-constrained minimization, but avoids duplication of variables and allows certain other economies in the implementation.

We give details of our approach by simplifying the notation and expressing the problem as follows:

$$\min_{z \in R^m} T_\lambda(z) := T(z) + \lambda \|z\|_1. \quad (30)$$

When  $T$  is convex (as in our application),  $z$  is optimal for (30) if and only if the following condition holds:

$$\nabla T(z) + \lambda v = 0, \quad (31)$$

where  $v$  is in the subdifferential of  $\|z\|_1$  (denoted by  $\partial\|z\|_1$ ), that is,

$$v_i \begin{cases} = -1 & \text{if } z_i < 0 \\ \in [-1, 1] & \text{if } z_i = 0 \\ = 1 & \text{if } z_i > 0 \end{cases} \quad (32)$$

A measure of near-optimality is given as follows:

$$\delta(z) = \min_{v \in \partial\|z\|_1} \|\nabla T(z) + \lambda v\|. \quad (33)$$

We have that  $\delta(z) = 0$  if and only if  $z$  is optimal.

The basic (first-order) step at iteration  $k$  is obtained by forming a simple model of the objective by expanding around the current iterate  $z^k$  as follows:

$$d^k = \arg \min_d T(z^k) + \nabla T(z^k)^T d + \frac{1}{2} \alpha_k d^T d + \lambda \|z^k + d\|_1, \quad (34)$$

where  $\alpha_k$  is a positive scalar (whose value is discussed below) and  $d^k$  is the proposed step. The subproblem (34) is separable in the components of  $d$  and therefore trivial to solve in closed form, in  $O(m)$  operations. We can examine the solution  $d^k$  to obtain an estimate of the active set as follows:

$$\mathcal{A}_k = \{i = 1, 2, \dots, m \mid (z^k + d^k)_i = 0\}. \quad (35)$$

We define the ‘‘inactive set’’ estimate  $\mathcal{I}_k$  to be the complement of the active set estimate, that is,

$$\mathcal{I}_k = \{1, 2, \dots, m\} \setminus \mathcal{A}_k.$$

If the step  $d^k$  computed from (34) does not yield a decrease in the objective function  $T_\lambda$ , we can increase  $\alpha_k$  and re-solve (34) to obtain a new  $d^k$ . This process can be repeated as needed. It can be shown that, provided  $z^k$  does not satisfy an optimality condition, the  $d^k$  obtained from (34) will yield  $T_\lambda(z^k + d^k) < T_\lambda(z^k)$  for  $\alpha_k$  sufficiently large.

We enhance the step by computing the restriction of the Hessian  $\nabla^2 T(z^k)$  to the set  $\mathcal{I}_k$  (which we denote by  $\nabla_{\mathcal{I}_k \mathcal{I}_k}^2 T(z^k)$ ) and then computing a Newton-like step in the  $\mathcal{I}_k$  components as follows:

$$(\nabla_{\mathcal{I}_k \mathcal{I}_k}^2 T(z^k) + \delta_k I) p_{\mathcal{I}_k}^k = -\nabla_{\mathcal{I}_k} T(z^k) - \lambda w_{\mathcal{I}_k}, \quad (36)$$

where  $\delta_k$  is a small damping parameter that goes to zero as  $z^k$  approaches the solution, and  $w_{\mathcal{I}_k}$  captures the gradient of the term  $\|z\|_1$  at the nonzero components of  $z^k + d^k$ .

Specifically,  $w_{\mathcal{I}_k}$  coincides with  $\partial\|z^k + d^k\|_1$  on the components  $i \in \mathcal{I}_k$ . If  $\delta_k$  were set to zero,  $p_{\mathcal{I}_k}^k$  would be the (exact) Newton step for the subspace defined by  $\mathcal{I}_k$ ; the use of a damping parameter ensures that the step is well defined even when the partial Hessian  $\nabla_{\mathcal{I}_k \mathcal{I}_k}^2 T(z^k)$  is singular or nearly singular, as happens with our problems. In our implementation, we choose

$$\delta_k = \min\left(\delta(z^k), \text{mean diagonal of } \nabla_{\mathcal{I}_k \mathcal{I}_k}^2 T(z^k)\right), \quad (37)$$

where  $\delta(z)$  is defined in (33).

Because of the special form of  $T(z)$  in our case (it is the function  $\mathcal{L}$  defined by (2) and (3)), the Hessian is not expensive to compute once the gradient is known. However, it is dense in general, so considerable savings can be made by evaluating and factoring this matrix on only a reduced subset of the variables, as we do in the scheme described above.

If the partial Newton step calculated above fails to produce a decrease in the objective function  $T_\lambda$ , we reduce its length by a factor  $\gamma_k$ , to the point where  $z_i^k + \gamma_k p_i^k$  has the same sign as  $z_i^k$  for all  $i \in \mathcal{I}_k$ . If this modified step also fails to decrease the objective  $T_\lambda$ , we try the first-order step calculated from (34), and take this step if it decreases  $T_\lambda$ . Otherwise, we increase the parameter  $\alpha_k$ , leave  $z^k$  unchanged, and proceed to the next iteration.

We summarize the algorithm as follows.

### Algorithm C.1

**given** *initial point*  $z^0$ , *initial damping*  $\alpha_0 > 0$ , *constants*  $\text{tol} > 0$  and  $\eta \in (0, 1)$ ;

**for**  $k = 0, 1, 2, \dots$

**if**  $\delta(z^k) < \text{tol}$

**stop** *with approximate solution*  $z^k$ ;

**end**

*Solve (34) for*  $d^k$ ; (*\* first-order step \**)

*Evaluate*  $\mathcal{A}_k$  and  $\mathcal{I}_k$ ;

*Compute*  $p_{\mathcal{I}_k}^k$  from (36); (*\* reduced Newton step \**)

*Set*  $z_{\mathcal{I}_k}^+ = z_{\mathcal{I}_k}^k + p_{\mathcal{I}_k}^k$  and  $z_{\mathcal{A}_k}^+ = 0$ ;

**if**  $T_\lambda(z^+) < \min(T_\lambda(z^k + d^k), T_\lambda(z^k))$  (*\* Newton step successful \**)

$z^{k+1} \leftarrow z^+$ ;

**else**

*Choose*  $\gamma_k$  as the largest positive number such that  $(z^k + \gamma_k p^k)_i z_i^k > 0$

    for all  $i$  with  $z_i^k \neq 0$ ; (*\* damp the Newton step \**)



```

Set  $z_{\mathcal{I}_k}^+ = z_{\mathcal{I}_k}^k + \gamma_k p_{\mathcal{I}_k}^k$  and  $z_{\mathcal{A}_k}^+ = 0$ ;
if  $T_\lambda(z^+) < \min(T_\lambda(z^k + d^k), T_\lambda(z^k))$  (* damped Newton step successful *)
     $z^{k+1} \leftarrow z^+$ ;
else if  $T_\lambda(z^k + d^k) < T_\lambda(z^k)$  (* first-order step successful; use it if Newton steps have
failed *)
     $z^{k+1} \leftarrow z^k + d^k$ ;
else (* unable to find a successful step *)
     $z^{k+1} \leftarrow z^k$ ;
end
end
(* increase or decrease  $\alpha$  depending on success of first-order step *)
if  $T_\lambda(z^k + d^k) < T_\lambda(z^k)$ 
     $\alpha_{k+1} \leftarrow \eta \alpha_k$ ; (* first-order step decreased  $T_\lambda$ , so decrease  $\alpha$  *)
else
     $\alpha_{k+1} \leftarrow \alpha_k / \eta$ ;
end
end

```

We have stated a simplified form of the algorithm above in which the second-order enhancement is calculated at *every* iteration. In practice, we compute it only when the number of components in  $\mathcal{I}_k$  is less than some threshold, taken as 500 in the experiments reported here.. We do so to avoid the expense of calculating and factoring a large reduced Hessian.

## D Results of Simulation Example 3

Table 7: Results of Simulation Example 3. In each row of a cell the first three numbers are the appearance frequencies of the important patterns and the last number is the appearance frequency of noise patterns.

$\rho_2 \setminus \rho_1$		0	0.2	0.5	0.7
0	LPS	96/100/100/54	98/100/100/46	100/100/100/43	100/100/100/44
	Logic	100/98/96/120	98/95/93/107	99/94/92/83	100/98/83/134
	SPLR	100/100/100/527	100/100/98/525	100/100/98/487	100/100/97/489
0.2	LPS	99/100/100/46	100/100/100/49	100/100/100/39	100/100/98/36
	Logic	99/97/94/96	100/99/87/94	100/100/88/73	100/99/86/117
	SPLR	100/100/94/517	100/99/96/530	100/97/95/495	100/100/96/485
0.5	LPS	99/100/99/47	99/100/100/51	100/100/99/51	100/100/98/46
	Logic	99/96/86/162	100/95/87/109	100/96/78/122	100/99/80/143
	SPLR	100/98/75/548	100/96/80/552	100/99/80/531	100/98/78/518
0.7	LPS	100/99/96/44	99/99/97/51	100/99/96/67	100/99/94/65
	Logic	100/83/70/195	100/88/69/167	100/85/70/153	100/89/74/126
	SPLR	100/91/51/580	100/85/49/594	100/81/52/584	100/72/55/570

## E Effect of Coding Flips

Proposition: Let  $f(x) = \mu + \sum c_{j_1 j_2 \dots j_r} B_{j_1 j_2 \dots j_r}(x)$  with all  $c_{j_1 j_2 \dots j_r}$  which appear in the sum strictly positive. If  $x_j \rightarrow 1 - x_j$  for  $j \in$  some subset of  $\{1, 2, \dots, p\}$  such that at least one  $x_j$  appears in  $f$ , then the resulting representation has at least one negative coefficient and at least as many terms as  $f$ . This follows from the

Lemma: Let  $g_k(x)$  be the function obtained from  $f$  by transforming  $x_j \rightarrow 1 - x_j, 1 \leq j \leq k$ . Then the coefficient of  $B_{j_1 j_2 \dots j_r}(x)$  in  $g_k(x)$  is

$$(-1)^{|\{j_1, \dots, j_r\} \cap \{1, \dots, k\}|} \sum_{\{j_1, \dots, j_r\} \subseteq T \subseteq \{j_1, \dots, j_r\} \cup \{1, \dots, k\}} c_T$$

where  $|\cdot|$  means number of entries.

## References

- [1] L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees*. Wadsworth, 1984.
- [2] K.-Y. Chan and W.-Y. Loh. Lotus: An algorithm for building accurate and comprehensible logistic regression trees. *Journal of Computational and Graphical Statistics*, 13:826–852, 2004.
- [3] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20:33–61, 1998.
- [4] P. Craven and G. Wahba. Smoothing noisy data with spline functions: estimating the correct degree of smoothing by the method of generalized cross-validation. *Numer. Math.*, 31:377–403, 1979.
- [5] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Ann. Statist.*, 32:407–499, 2004.
- [6] H. Cordell et al. Genetic analysis workshop 15: gene expression analysis and approaches to detecting multiple functional loci. *BMC Proceedings*, 1(Suppl 1):S1,1–4, 2007.
- [7] W. Thompson et al. Rheumatoid arthritis association at 6q23. *Nature Genetics*, Nov 4, 2007. doi:10.1038/ng.2007.32.
- [8] P. Galan, F. Viteri, and S. Bertrais et al. Serum concentrations of beta-carotene, vitamins C and E, zinc and selenium are influenced by sex, age, diet, smoking status, alcohol consumption and corpulence in a general French adult population. *Eur. J. Clin. Nutr.*, 59:1181–90, 2005.
- [9] F. Gao, G. Wahba, R. Klein, and B. Klein. Smoothing spline ANOVA for multivariate Bernoulli observations, with applications to ophthalmology data, with discussion. *J. Amer. Statist. Assoc.*, 96:127–160, 2001.
- [10] E. George. The variable selection problem. *J. Amer. Statist. Assoc.*, 95:1304–1308, 2000.

- [11] G. Golub, M. Heath, and G. Wahba. Generalized cross validation as a method for choosing a good ridge parameter. *Technometrics*, 21:215–224, 1979.
- [12] S. Gunn and J. Kandola. Structural modelling with sparse kernels. *Machine Learning*, 48:137–163, 2002.
- [13] D. Haughton. On the choice of a model to fit data from an exponential family. *Ann. Statist.*, 16:342–355, 1988.
- [14] R. Klein, B. E. K. Klein, K. Linton, and D. DeMets. The Beaver Dam eye study: Visual acuity. *Ophthalmology*, 98:1310–1315, 1991.
- [15] K. Knight and W. Fu. Asymptotics for LASSO-type estimators. *Ann. Statist*, 28:1356–1378, 2000.
- [16] K. Lee, B. Klein, and R. Klein. Changes in refractive error over a 5-year interval in the beaver dam eye study. *Investigative Ophthalmology & Visual Science*, 40:1645–1649, 1999.
- [17] K. Lee, B. Klein, R. Klein, and T. Wong. Changes in refraction over 10 years in an adult population: the Beaver Dam Eye Study. *Investigative Ophthalmology Visual Science*, 43:2566–71, 2002.
- [18] Y. Lee, Y. Kim, S. Lee, and J. Koo. Structured multicategory support vector machines with analysis of variance decomposition. *Biometrika*, 93:555–571, 2006.
- [19] C. Leng, Y. Lin, and G. Wahba. A note on the LASSO and related procedures in model selection. *Statistica Sinica*, 16:1273–1284, 2006.
- [20] K. C. Li. Asymptotic optimality of  $C_L$  and generalized cross validation in ridge regression with application to spline smoothing. *Ann. Statist.*, 14:1101–1112, 1986.
- [21] X. Lin. Smoothing spline analysis of variance for polychotomous response data. Technical Report 1003, PhD thesis, Department of Statistics, University of Wisconsin, Madison WI, 1998. Available via G. Wahba’s website.
- [22] M. Miller, G. Lind, N. Li, and S. Jang. Genetic analysis workshop 15: Simulation of a complex genetic model for rheumatoid arthritis in nuclear families including a dense snp

- map with linkage disequilibrium between marker loci and trait loci. *BMC Proceedings*, 1(Suppl 1):S4, 1–7, 2007.
- [23] M. Osborne, B. Presnell, and B. Turlach. On the LASSO and its dual. *J. Comp. Graph. Stat.*, 9:319–337, 2000.
- [24] M. Park and T. Hastie. Penalized logistic regression for detecting gene interactions. Technical Report 00-25, Department of Statistics, Stanford University, 2006.
- [25] I. Ruczinski, C. Kooperberg, and M. LeBlanc. Logic regression. *J. Computational and Graphical Statistics*, 12:475–511, 2003.
- [26] D. Schwartz, S. Szymczak, A. Ziegler, and R. Konig. Picking single-nucleotide polymorphisms in forests. *BMC Proceedings*, 1(Suppl 1):S59, 1–5, 2007.
- [27] B. Seet, T. Wong, D. Tan, S. Saw, V. Balakrishnan, L. Lee, and A. Lim. Myopia in Singapore: taking a public health approach. *Br. J. Ophthalmol.*, 85:521–526, 2001.
- [28] W. Shi, K. Lee, and G. Wahba. Detecting disease-causing genes by LASSO-patternsearch algorithm. *BMC Proceedings*, 1(Suppl 1):S60, 1–5, 2007.
- [29] W. Shi, G. Wahba, S. Wright, K. Lee, R. Klein, and B. Klein. LASSO-Patternsearch algorithm with application to ophthalmology data. Technical Report 1131, Department of Statistics, University of Wisconsin, Madison WI, 2006. Under revision.
- [30] R. Tibshirani. Regression shrinkage and selection via the LASSO. *J. Roy. Stat. Soc, B*, 58:267–288, 1996.
- [31] G. Wahba. Soft and hard classification by reproducing kernel Hilbert space methods. *Proceedings of the National Academy of Sciences*, 99:16524–16530, 2002.
- [32] G. Wahba and S. Wold. A completely automatic French curve. *Commun. Stat.*, 4:1–17, 1975.
- [33] J. Whittaker. *Graphical Models in Applied Mathematical Multivariate Statistics*. Wiley, 1990.
- [34] D. Xiang and G. Wahba. A generalized approximate cross validation for smoothing splines with non-Gaussian data. *Statistica Sinica*, 6:675–692, 1996.

- [35] H. Zhang and Y. Lin. Component selection and smoothing for nonparametric regression in exponential families. *Statistica Sinica*, 16:1021–1042, 2006.
- [36] H. Zhang, G. Wahba, Y. Lin, M. Voelker, M. Ferris, R. Klein, and B. Klein. Variable selection and model building via likelihood basis pursuit. *J. Amer. Statist. Assoc.*, 99:659–672, 2004.
- [37] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B*, 67:301–320, 2005.