

Convex Relaxations for Subset Selection

Francis Bach*, Selin Damla Ahipařaođlu†, Alexandre d'Aspremont‡

June 17, 2010

Abstract

We use convex relaxation techniques to produce lower bounds on the optimal value of subset selection problems and generate good approximate solutions. We then explicitly bound the quality of these relaxations by studying the approximation ratio of sparse eigenvalue relaxations. Our results are used to improve the performance of branch-and-bound algorithms to produce exact solutions to subset selection problems.

1 Introduction

We focus here on the subset selection problem., i.e., solving least squares regressions while constraining the number of nonzero regression variables to be less than a certain target. This problem is often called feature selection or sparse least-squares. Its combinatorial nature makes subset selection intractable. Several techniques have been derived to produce good approximate solutions however, using for example greedy algorithms or sparsity inducing penalties.

Given a design matrix $X \in \mathbf{R}^{n \times p}$ and a response vector $y \in \mathbf{R}^n$, we consider the following subset selection problem

$$\begin{aligned} & \text{minimize} && \|y - Xw\|_2^2 \\ & \text{subject to} && \mathbf{Card}(w) \leq k, \end{aligned} \tag{1}$$

in the variable $w \in \mathbf{R}^p$, where k is a parameter controlling sparsity. It was shown in Natarajan (1995) that while (1) is NP-Hard, simple greedy algorithms can efficiently produce good approximate solutions. Subset selection can also be understood as ℓ_0 norm constrained regression (or approximation) and a very large body of works focused on replacing the combinatorial ℓ_0 norm with a convex ℓ_1 norm constraint, with ℓ_1 norm regression usually known as LASSO Tibshirani (1996). Explicit variable selection consistency results have been derived in certain regimes (see e.g. Meinshausen et al. (2007)), and recent results Donoho and Tanner (2005); Candes and Tao (2007) have shown that under certain conditions on the design matrix X , the solutions of the ℓ_1 problem coincided with that of the ℓ_0 problem. Several authors have attacked the ℓ_0 problem directly, with Narendra and Fukunaga (1977); Hand (1981); Furnival and Wilson Jr (2000); Moghaddam et al. (2008) using branch-and-bound techniques to produce exact solutions to problem (1), with Moghaddam et al. (2008) in particular using interlacing properties of eigenvalues to

*INRIA-WILLOW Project-Team, Laboratoire d'Informatique de l'Ecole Normale Supérieure (CNRS/ENS/INRIA UMR 8548), 23, avenue d'Italie, 75214 Paris, France. francis.bach@mines.org

†ORFE, Princeton University, Princeton, NJ 08544. sahipasa@princeton.edu

‡ORFE, Princeton University, Princeton, NJ 08544. aspremon@princeton.edu

speedup branch-and-bound methods. Solving the ℓ_0 problem in (1) even for small values of p has direct applications in image denoising Elad and Aharon (2006); Mairal et al. (2008).

All the algorithms listed above produce good approximate solutions, hence *upper bounds* on the optimal value of the subset selection problem (1). Our first contribution here is to use convex relaxation techniques to produce *lower bounds* on the optimal value of (1). In particular, this result allows us to bound the suboptimality of approximate solutions and improve the performance of branch-and-bound algorithms for subset selection. We also use randomization techniques to generate good solutions to (1), often improving on solutions produced by greedy or LASSO algorithms. Our next main contribution is to derive approximation bounds on the performance of the sparse eigenvalue relaxation/randomization algorithm. Finally, we test our algorithms on various subset selection problems and show that the lower bound derived here considerably reduces the number of branches required to produce an optimal solution to (1).

The paper is organized as follows. In Section 2 we show how to produce lower bounds on the optimal value of problem (1) using relaxation bounds on sparse eigenvalues. In Section 3, we describe greedy, randomization and branch-and-bound algorithms to generate good approximate solutions w to (1) using the product of the relaxation. In Section 4 we produce a bound on the approximation ratio of sparse eigenvalues, thus bounding the quality of the approximation of the subset selection bounds derived in Section 2. Section 5 shows how to efficiently solve our semidefinite relaxation using first-order methods. Finally, Section 6 presents some numerical experiments.

Notations

Given matrices $X, Y \in \mathbf{R}^{n \times p}$, we write $X \circ Y$ their Schur (componentwise) product, while $\lambda_{\max}(X)$ is the leading eigenvalue of X , $\|X\|_1$ the sum of absolute values of the coefficients in X and X_i is the i^{th} column of X . We let \mathbf{S}_p be the set of symmetric matrices, and for $Y \in \mathbf{S}_p$, we write $\mathbf{diag}(Y) \in \mathbf{R}^p$ its diagonal. When $y \in \mathbf{R}^p$, $\mathbf{diag}(y) \in \mathbf{S}_p$ denotes de diagonal matrix with diagonal coefficients equal to the coefficients of y , while $\mathbf{Card}(y)$ is the number of nonzero coefficients in y .

2 Relaxation & Lower Bounds

Following d'Aspremont et al. (2008) for example, we first recall how solving problem (1) is equivalent to computing sparse eigenvalues of a matrix formed using X and y . We let $\psi(k)$ be the optimal value of the subset selection problem (1), with

$$\psi(k) = \min_{\substack{\|y - Xw\|_2^2 \\ \mathbf{Card}(w) \leq k,}} \quad (2)$$

in the variable $w \in \mathbf{R}^p$, where k is again a parameter controlling sparsity. We can rewrite this

$$\begin{aligned} \psi(k) &= \min_{\substack{\mathbf{1}^T u \leq k \\ u \in \{0,1\}^p}} \min_{\substack{\mathbf{Card}(w) \leq k \\ w \in \mathbf{R}^p}} \|y - X \mathbf{diag}(u)w\|_2^2 \\ &= \min_{\substack{\mathbf{1}^T u \leq k \\ u \in \{0,1\}^p}} \min_{w \in \mathbf{R}^p} \|y - X \mathbf{diag}(u)w\|_2^2 \\ &= \min_{\substack{\mathbf{1}^T u \leq k \\ u \in \{0,1\}^p}} \min_{\|w\|_2=1} \min_{\nu \in \mathbf{R}} \|y - X \mathbf{diag}(u)\nu w\|_2^2, \end{aligned}$$

which, after minimizing explicitly over ν , becomes

$$\min_{\substack{\mathbf{1}^T u \leq k \\ u \in \{0,1\}^p}} \min_{\|w\|_2=1} y^T y - \frac{(y^T X \mathbf{diag}(u)w)^2}{w^T \mathbf{diag}(u)X^T X \mathbf{diag}(u)w}.$$

This means that $\psi(k) \leq y^T y - \rho$ if and only if

$$\max_{\substack{\mathbf{1}^T u \leq k \\ u \in \{0,1\}^p}} \max_{\|w\|_2=1} \frac{(y^T X \mathbf{diag}(u)w)^2}{w^T \mathbf{diag}(u)X^T X \mathbf{diag}(u)w} \leq \rho.$$

We can rewrite this condition

$$wX^T(yy^T - \rho\mathbf{I})Xw \leq 0, \quad \text{when } \|w\|_2 \leq 1, \mathbf{Card}(w) \leq k$$

which is equivalent to

$$\lambda_{\max}^k(X^T(yy^T - \rho\mathbf{I})X) \leq 0. \quad (3)$$

Here, $\lambda_{\max}^k(A)$ is the k sparse maximum eigenvalue of a matrix $A \in \mathbf{S}_p$, defined as

$$\lambda_{\max}^k(A) = \max_{\substack{x^T Ax \\ \text{s. t. } \|x\| = 1, \mathbf{Card}(x) \leq k}} \quad (4)$$

in the variable $x \in \mathbf{R}^p$. Relaxation bounds for sparse eigenvalues $\lambda_{\max}^k(A)$ were derived in d'Aspremont et al. (2007, 2008), with the bound in d'Aspremont et al. (2007) written

$$\lambda_{\max}^k(A) \leq \begin{array}{ll} \text{maximize} & \mathbf{Tr} AZ \\ \text{subject to} & \|Z\|_1 \leq k \\ & \mathbf{Tr}(Z) = 1, Z \succeq 0 \end{array} \quad (5)$$

in the variable $Z \in \mathbf{S}_p$. We can summarize the above derivation in the following proposition.

Proposition 1 *Given a design matrix $X \in \mathbf{R}^{n \times p}$ and a response vector $y \in \mathbf{R}^n$, consider the following subset selection problem*

$$\psi(k) = \begin{array}{ll} \text{min.} & \|y - Xw\|_2^2 \\ \text{s.t.} & \mathbf{Card}(w) \leq k, \end{array}$$

then

$$\psi(k) \geq y^T y - \rho \quad \text{if and only if } \lambda_{\max}^k(X^T(yy^T - \rho\mathbf{I})X) \leq 0,$$

where $\lambda_{\max}^k(\cdot)$ is the sparse maximum eigenvalue function defined in (4).

3 Approximate Solutions

The relaxation detailed in (5) produces a *lower bound* on the objective value of problem (1). In this section, we describe how to use the solution of this relaxation to produce good approximate solution vectors w to problem (1), hence produce *upper bounds* on the solution value. We first describe greedy algorithms which can be used to solve problem (1) independently, or to improve solutions extracted from convex relaxations.

3.1 Greedy methods

To simplify notations here, we first define the following function, which computes the solution value of problem (1) *given* the support of the solution vector $w \in \mathbf{R}^p$. Let $I \subset [1, p]$ be a index subset such that $w_i = 0$ if $i \notin I$, we write I^c its complement in $[1, p]$ and let

$$\mu(I) = \min_{w_{I^c}=0} \|y - Xw\|_2^2 \tag{6}$$

in the variable $w \in \mathbf{R}^p$. Note that while computing the optimal value of problem (1) is NP-Hard, computing $\mu(I)$ in (6) is equivalent to forming a QR decomposition of the matrix $X_I \in \mathbf{R}^{n \times k}$ where k is the cardinality of the support set I .

We can greedily construct approximate solutions to (6) by scanning variables at each iteration to increase (or decrease) the size of the support as in the *forward greedy* algorithm is detailed in Algorithm 1. The *backward greedy algorithm* is similar but starts from the full support $[1, p]$ and progressively removes points.

Algorithm 1 Forward Greedy Algorithm.

Input: $X \in \mathbf{R}^{n \times p}$, $y \in \mathbf{R}^n$, target cardinality k^{target} .

- 1: Initialization: $I_0 = \emptyset$.
- 2: **for** $i = 1$ to k^{target} **do**
- 3: Compute $i_k = \operatorname{argmax}_{i \notin I_{k-1}} \mu(I_{k-1} \cup \{i\})$
- 4: Set $I_k = I_{k-1} \cup \{i_k\}$ and compute w_k as the minimizer of $\mu(I_k)$ in (6).
- 5: **end for**

Output: Support sets I_k for w in problem (1), with $k = 1, \dots, k^{\text{target}}$.

3.2 Randomization

As in the MAXCUT relaxation by Goemans and Williamson (1995) for example, we can use the matrix solution to the relaxation in (5) to generate good approximate solutions to problem (1). The solution matrix Z in (5) can be understood as a covariance matrix, and we use it to generate Gaussian vectors $z \sim \mathcal{N}(0, Z)$. The k indices corresponding to the k largest magnitude coefficients of the sample vectors z then provide support sets I corresponding to nonzero coefficients in w . Given these support sets, one then solves for $\mu(I)$ in (6) to get upper bounds on the optimal value of (1) and approximate solution vectors w .

In the next section, we will also consider another much simpler randomization procedure whose performance can be completely characterized. This second procedure does not require solving relaxation (5), but simply computing a leading eigenvector x of the matrix A in (4). Good approximate solutions $z \in \mathbf{R}^p$ to problem (4) are then randomly sampled with $z_i = 1/\sqrt{k}$ with probability $p_i = k|x_i|/\|x\|_1$ and $z_i = 0$ otherwise. We then prune z using a few backward greedy step, whenever $\mathbf{Card}(z) > k$. While the complexity of this procedure is much lower than that of the full greedy algorithm, we will see in the next section that it produces solutions of comparable quality.

3.3 Branch-and-bound algorithm

As in Furnival and Wilson Jr (2000); Moghaddam et al. (2008), we can develop a branch-and-bound algorithm for finding optimal solutions to (1). Suppose we are looking for a vector in \mathbf{R}^p with at

most k non-zero components, we need to enumerate at most $\binom{p}{k}$ subsets to find the best one. We start by dividing all possible subsets into two branches, one containing the first variable and one which does not. We further branch each of these branches into two, one containing the second variable and one not, etc. At each node of the search tree, we have a subproblem that excludes certain variables (depending on branching decisions made so far). For each subproblem, we generate lower bounds using Proposition 1 by solving relaxation (5), and upper bounds when there are exactly k variables left on the branch. We also generate upper bounds by applying a combination of the greedy algorithms and randomization techniques described above to the solutions of the relaxed problems. Obviously, we fathom a node whose lower bound exceeds the best upper bound since the branches diverging from this node cannot contain a better solution than the best solution found so far.

4 Tightness

The sparse eigenvalue problem in (4) is closely connected to the k -Dense-Subgraph problem described in Kortsarz and Peleg (1993); Feige et al. (2001); Feige and Langberg (2001) for example. The k -Dense-Subgraph problem seeks to find a principal submatrix of A of dimension k with largest coefficient sum. This is written

$$\max_{\mathbf{1}^T u \leq k} u^T A u$$

in the variable $u \in \{0, 1\}^p$. On the other hand, the problem of computing a sparse maximum eigenvalue can be written

$$\lambda_{\max}^k(A) = \max_{\mathbf{1}^T u \leq k} \max_{\|x\|=1} u^T (A \circ xx^T) u$$

in the variables $x \in \mathbf{R}^p$, $u \in \{0, 1\}^p$. We thus observe that computing sparse eigenvalues means solving a k -Dense-Subgraph problem over the result of an inner eigenvalue problem in x . Below, we first recall an approximation result on the backward greedy algorithm used in Moghaddam et al. (2008), which applies to *positive semidefinite* matrices A .

Proposition 2 *Let $A \in \mathbf{S}_p$, with $A \succeq 0$ and $k > 0$ and suppose $\mathbf{diag}(A) \geq 0$. We have*

$$\frac{k}{p} \lambda_{\max}(A) \leq \lambda_{\max}^k(A) \leq \lambda_{\max}(A) \tag{7}$$

where $\lambda_{\max}^k(A)$ is the optimal value of problem (4).

Proof. From (Horn and Johnson, 1985, §4.3.14), when $A \succeq 0$, we have

$$\lambda_{\max}^i(A) \geq \frac{i}{i+1} \lambda_{\max}^{i+1}(A)$$

for any $i \in [1, p-1]$. A simple recursion then gives the desired result. ■

When A is not positive semidefinite, we can adapt results from Feige and Seltser (1997) to show

$$\frac{k(k-1)}{p(p-1)} \lambda_{\max}(A) \leq \lambda_{\max}^k(A) \leq \lambda_{\max}(A). \tag{8}$$

When the coefficients of A are nonnegative, we can obtain approximation bounds for basic randomization techniques similar to those developed in Feige and Seltser (1997) for the k -Dense-Subgraph problem. The approximation ratio in this case also decreases as k/p , which shows that the randomization algorithm has a performance comparable to that of the backward greedy method, while being significantly cheaper on large scale problems.

Proposition 3 *Let $A \in \mathbf{S}_p$ such that $A_{ij} \geq 0$, $i, j = 1, \dots, p$ and $k > 1$. The sparse eigenvalue problem defined in (4) was written*

$$\begin{aligned} \lambda_{\max}^k(A) = & \max. \quad x^T A x \\ \text{s. t.} \quad & \|x\|_2 \leq 1, \mathbf{Card}(x) \leq k \end{aligned}$$

in the variable $x \in \mathbf{R}^p$. We then have

$$\frac{k}{p} \mu(k, p) \lambda_{\max}(A) \leq \lambda_{\max}^k(A) \leq SDP_k(A) \leq \lambda_{\max}(A) \quad (9)$$

where

$$\mu(k, p) = \left(1 - \frac{2}{k^{1/3}}\right) \left(1 - \frac{p^2}{k^2} e^{-p^{1/9}/3}\right) \xrightarrow{p \rightarrow \infty} 1 \quad (10)$$

whenever $k \geq p^{1/3}$, for p sufficiently large, where $SDP_k(A)$ is the optimal value of (5).

Proof. To maintain the parallel with Feige and Seltser (1997), we write $Z = xx^T$, where x is a leading eigenvector of A . The matrix Z then satisfies $\mathbf{Tr} Z = 1$ and $Z \succeq 0$. The upper bound in (9) follows directly from d'Aspremont et al. (2007) and we focus here on the lower bound. We randomly sample vectors $z \in \mathbf{R}^p$ such that

$$z_i = \begin{cases} 1/\sqrt{k} & \text{with probability } p_i = k\sqrt{Z_{ii}}/S, \\ 0 & \text{otherwise.} \end{cases}$$

where $S = \sum_{i=1}^p \sqrt{Z_{ii}}$. We then have

$$\begin{aligned} \mathbf{E}[z^T A z] &= \mathbf{Tr}(A \mathbf{E}[z z^T]) \\ &= \sum_{i,j=1}^p k A_{ij} \sqrt{Z_{ii} Z_{jj}} / S^2 \\ &\geq \sum_{i,j=1}^p k A_{ij} \sqrt{Z_{ii} Z_{jj}} / p \\ &\geq \frac{k}{p} \mathbf{Tr}(A Z) \end{aligned}$$

where the first inequality uses $\mathbf{Tr} Z = 1$ and the last (Cauchy) inequality follows from the fact that $Z \succeq 0$ and $A \geq 0$. Now, let $q = \mathbf{Prob}[z^T A z \leq \mathbf{E}[z^T A z]/\beta]$ for some $\beta \geq 1$, we have

$$\mathbf{E}[z^T A z] \leq q \frac{\mathbf{E}[z^T A z]}{\beta} + (1 - q) \frac{\mathbf{1}^T A \mathbf{1}}{k}$$

which means

$$q \leq 1 - \frac{\beta - 1}{\beta \mathbf{1}^T A \mathbf{1} / k \mathbf{E}[z^T A z] - 1}.$$

because $\mathbf{1}^T A \mathbf{1} / k \geq \mathbf{E}[z^T A z]$. Now, using Chernoff's inequality as in (Feige and Seltser, 1997, Lem. 4.1) produces

$$\mathbf{Prob} [\mathbf{Card}(z) - \mathbf{1}^T p \geq t \mathbf{1}^T p] \leq e^{-\frac{t^2 \mathbf{1}^T p}{3}},$$

so, as in (Feige and Seltser, 1997, Th. 4.1), when $k \geq p^{1/3}$

$$\mathbf{Prob} \left[\mathbf{Card}(z) \geq k \left(1 + k^{-1/3} \right) \right] \leq e^{-p^{1/9}/3}.$$

We have

$$\frac{\mathbf{1}^T A \mathbf{1}}{k \mathbf{E}[z^T A z]} \leq \frac{p \mathbf{1}^T A \mathbf{1}}{k^2 \mathbf{Tr} AZ} \leq \frac{p^2 \lambda_{\max}(A)}{k^2 \mathbf{Tr} AZ} = \frac{p^2}{k^2}$$

which follows from $A \geq 0$, $\mathbf{Tr} AZ = \lambda_{\max}(A)$, $\mathbf{Tr} Z = 1$ with $Z \succeq 0$. When $k \geq p^{1/3}$ and p is large enough so that $p^2 e^{-p^{1/9}/3} / k^2 < 1$, we can enforce

$$\beta > \frac{1}{1 - p^2 e^{-p^{1/9}/3} / k^2}$$

and thus get

$$\beta > \frac{e^{p^{1/9}/3} - 1}{e^{p^{1/9}/3} - \mathbf{1}^T A \mathbf{1} / k \mathbf{E}[z^T A z]},$$

which means, using the bound on q derived above,

$$1 - q \geq \frac{\beta - 1}{\beta \mathbf{1}^T A \mathbf{1} / k \mathbf{E}[z^T A z] - 1} > e^{-\frac{t^2 \mathbf{1}^T p}{3}},$$

which, combined with the deviation bounds detailed above, yields

$$\mathbf{Prob}[z^T A z \geq \mathbf{E}[z^T A z] / \beta] = 1 - q > e^{-p^{1/9}/3} \geq \mathbf{Prob} \left[\mathbf{Card}(z) \geq k \left(1 + k^{-1/3} \right) \right].$$

This shows that by sampling enough points z , we can generate a vector $z_0 \in \mathbf{R}^p$ such that

$$z_0^T A z_0 \geq \frac{k}{\beta p} \mathbf{Tr}(AZ) \quad \text{and} \quad \mathbf{Card}(z_0) \leq k \left(1 + k^{-1/3} \right)$$

If we remove at most $k^{2/3}$ variables from z_0 using the backward greedy algorithm described in the previous section, (8) shows that we loose at most a factor

$$\frac{k(k-1)}{(k+k^{2/3})(k+k^{2/3}-1)} = 1 - \frac{2}{k^{1/3}} + \frac{3}{k^{2/3}} + o\left(\frac{1}{k^{2/3}}\right)$$

and, when p is large enough, we obtain a point z_k such that

$$z_k^T A z_k \geq \frac{k}{p} \left(1 - \frac{2}{k^{1/3}} \right) \frac{\mathbf{Tr}(AZ)}{\beta}, \quad \|z_k\|_2 \leq 1 \quad \text{and} \quad \mathbf{Card}(z_k) \leq k,$$

which means that z_k is a feasible point of problem (4), and yields the desired result. \blacksquare

Note that the randomization procedure detailed in the proof above is simpler than the one we used in Section 3.2, producing bounds on the performance of the later one is unfortunately much harder. We can directly extend this last proposition to problems where A has negative coefficients, but the bound is not proportional in this case.

Proposition 4 Let $A \in \mathbf{S}_p$ and $k > 1$. We have

$$\min \left\{ 0, \min_{i,j=1,\dots,p} A_{ij} \right\} k + \frac{k}{p} \mu(k,p) \lambda_{\max}(A) \leq \lambda_{\max}^k(A) \leq SDP_k(A) \leq \lambda_{\max}(A), \quad (11)$$

where $SDP_k(A)$ is the optimal value of (5) and $\mu(k,p)$ is defined in (10), whenever $k \geq p^{1/3}$ and p is sufficiently large.

Proof. The function $\lambda_{\max}^k(\cdot)$ defined in (4) is convex as a pointwise maximum of affine functions. This implies

$$\lambda_{\max}^k(A) \geq \lambda_{\max}^k(A - \min_{ij} A_{ij} \mathbf{1}\mathbf{1}^T) + \min_{ij} A_{ij} (\mathbf{1}^T x)^2$$

for some vector x satisfying $\|x\|_2 = 1$ and $\mathbf{Card}(x) \leq k$. The matrix $A - \min_{ij} A_{ij} \mathbf{1}\mathbf{1}^T$ is nonnegative and Proposition 3 shows that

$$\lambda_{\max}^k(A - \min_{ij} A_{ij} \mathbf{1}\mathbf{1}^T) \geq \frac{k}{p} \mu(k,p) \lambda_{\max}(A - \min_{ij} A_{ij} \mathbf{1}\mathbf{1}^T).$$

We then get

$$\begin{aligned} & \frac{k}{p} \mu(k,p) \lambda_{\max}(A - \min_{ij} A_{ij} \mathbf{1}\mathbf{1}^T) + \min_{ij} A_{ij} (\mathbf{1}^T x)^2 \\ & \geq \frac{k}{p} \mu(k,p) (\lambda_{\max}(A) - \min_{ij} A_{ij} (\mathbf{1}^T y)^2) + \min_{ij} A_{ij} (\mathbf{1}^T x)^2 \\ & \geq \frac{k}{p} \mu(k,p) \lambda_{\max}(A) + \min_{ij} A_{ij} (\mathbf{1}^T x)^2 \end{aligned}$$

when $\min_{ij} A_{ij} < 0$, which follows from the convexity of $\lambda_{\max}(\cdot)$, where y is a leading eigenvector of $A - \min_{ij} A_{ij} \mathbf{1}\mathbf{1}^T$. We conclude using $(\mathbf{1}^T x)^2 \leq \|x\|_1^2 \leq \mathbf{Card}(x) \|x\|_2^2 = k$. ■

5 Convex Minimization Algorithm

The relaxation in Section 2 meant solving

$$\begin{aligned} & \text{maximize} && \mathbf{Tr} MZ \\ & \text{subject to} && \|Z\|_1 \leq k \\ & && \mathbf{Tr}(Z) = 1, Z \succeq 0 \end{aligned}$$

in the variable $Z \in \mathbf{S}_p$, where $M \in \mathbf{S}_p$ was formed as $M = X^T (yy^T - \rho \mathbf{I}) X$. We compute the dual of this problem by first writing it in a saddle-point format.

$$\min_{\lambda \geq 0} \max_{\substack{\mathbf{Tr}(Z)=1 \\ Z \succeq 0}} \mathbf{Tr} MZ + \lambda(k - \|Z\|_1)$$

which is also

$$\min_{\lambda \geq 0} \max_{\substack{\mathbf{Tr}(Z)=1 \\ Z \succeq 0}} \min_{\|Y\|_\infty \leq 1} \mathbf{Tr} Z(M + \lambda Y) + k\lambda$$

in the variables $Z, Y \in \mathbf{S}_p$. We can rewrite this as

$$\min_{Y \in \mathbf{S}_p} \max_{\substack{\text{Tr}(X)=1 \\ X \succeq 0}} \text{Tr} X(M + Y) + k\|Y\|_\infty$$

which is equivalent to

$$\min \lambda_{\max}(M + Y) + k\|Y\|_\infty \tag{12}$$

in the variable $Y \in \mathbf{S}_p$. This is a maximum eigenvalue minimization problem and can be solved efficiently using for example smooth first-order algorithms as in Nesterov (2003). Given an a priori bound on suboptimality, the total complexity of obtaining a solution up to accuracy ϵ then grows as

$$O\left(\frac{kn^3\sqrt{\log n}}{\epsilon}\right).$$

Given an approximate solution $Y \in \mathbf{S}_p$ to the dual, we can reconstruct a corresponding primal solution Z by first solving

$$Z = \operatorname{argmax}_{\substack{\text{Tr}(Z)=1 \\ Z \succeq 0}} \text{Tr} Z(M + Y)$$

and checking if $\|Z\|_1 \leq k$ (this last condition will always be satisfied if Y is optimal).

6 Numerical Results

Table 1 presents numerical experiments using branch-and-bound on a set of small artificial problems. We generate normally distributed matrices $X \in \mathbf{R}^{n \times p}$, a random sparse vector w whose cardinality is at most k , and a righthand side vector y , which is equal to $Xw + \epsilon$, where $\epsilon \in \mathbf{R}^n$ is noise. The last four columns are related to the performance of the B&B algorithm: the first gives the smallest number of nodes visited by the algorithm, the second provides the average number of nodes visited over all instances, the third shows the number of nodes in the complete enumeration tree while the fourth lists the average speedup. These results suggest that the lower bound obtained in this paper is effective in fathoming a significant number of nodes in the search tree. Out of these 160 small test instances, the forward greedy algorithm found the optimal solution for 105 problems, whereas the randomization algorithm followed by a greedy improvement step (which will be referred as the enhanced randomization algorithm from now on) was able to find the optimal solution for 113 problems. Unfortunately, the authors of Moghaddam et al. (2008) did not release a software package and the “leaps and bounds” package released by the authors of Furnival and Wilson Jr (2000) does not output the number of nodes it visits so direct comparisons were not possible.

Table 1: Number of nodes visited by the branch-and-bound algorithm.

p	n	k	No. instances	B&B (Best)	B&B (Average)	$\binom{p}{k}$	Speedup (Avg.)
20	10	2	100	35	194	380	2
30	15	3	50	330	4 799	24 360	5
40	20	4	10	42 236	98 236	2 193 360	22
50	25	4	2	71 552	96 734	5 527 200	57

On larger instances where $p = 100$ and $n = 50$, the cardinality of w was set to 2 and 4. Figure 1 plots lower bounds (Low. Bnd.) on (1) generated by solving relaxation (5), the coarse solution points (Primal) extracted from the matrix Z solving (5), the solutions (Greedy) obtained by the forward greedy algorithm, the LARS algorithm Efron et al. (2004), and the enhanced solutions (Rand) obtained by applying the randomization algorithm detailed in Section 3.2 to the matrix Z solving (5). We observe that around the true cardinality of w used in generating the problem instances, the enhanced relaxation sometimes outperforms both the forward greedy algorithm and LARS and always performs at least as good as the best of these two methods.

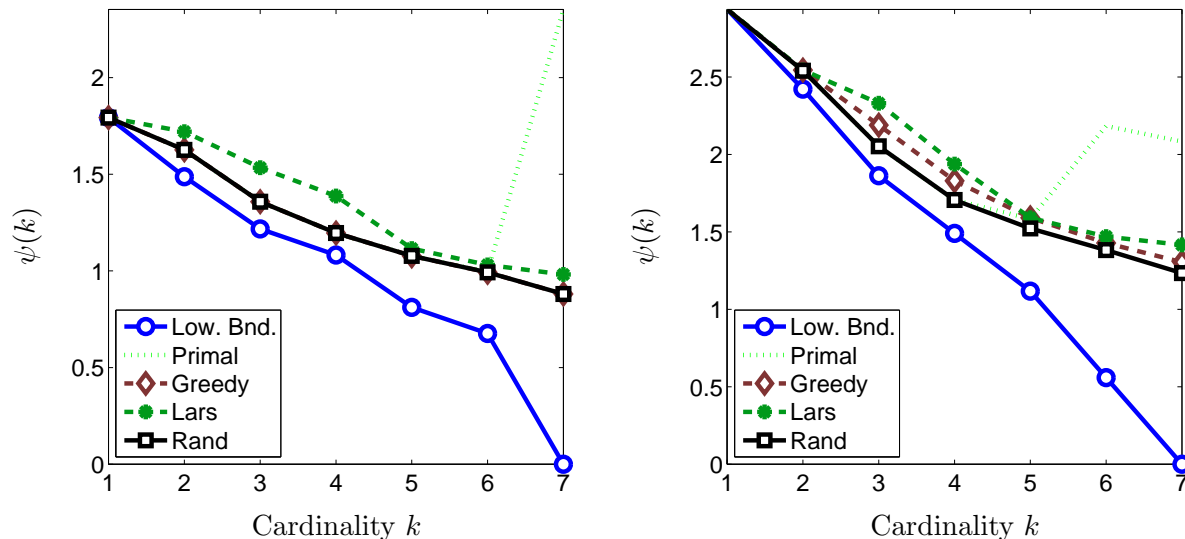


Figure 1: Lower bounds and objective values in (1) for various algorithms, versus target cardinality. The true cardinality is 2 (left) and 4 (right).

More realistic data sets were generated with an image compression setting in mind. $X \in \mathbf{R}^{n \times p}$ is now an overcomplete dictionary of Gabor wavelets, and y is an image patch of size $r \times r$ obtained from an actual image. We set $r = 10$ for all the experiments. We first solve this batch of problems (for $p = 24$ and $n = 16$) with the B&B algorithm where the target cardinality is either 2, 3, or 4. We then compare the performance of the forward greedy algorithm and the enhanced randomization algorithm of Section 3.2. Table 2 shows that the modified randomization algorithm finds the optimal solution in most cases.

Table 2: Number of instanced solved by greedy and randomization algorithms on image data.

Dimensions			Greedy			Randomization	
p	n	k	No. instances	No. solved	Max. Rel. Gap	No. Solved	Max. Rel. Gap
24	16	2	10	9	0.22	9	0.90
24	16	3	10	8	0.70	9	0.16
24	16	4	10	8	0.94	9	0.31

Most of our experiments so far were focused on finding *exact* solutions to small instances of problem (1). We also tested the numerical complexity of our methods on larger problems for which we only sought good upper and lower bounds. Computing times for solving relaxation (5) on increasingly large Gaussian random problems (generated as above) are reported in Table 3.

Table 3: CPU time versus problem size.

Problem size p	CPU time
100	0 h 00 m 07 s
250	0 h 01 m 32 s
500	0 h 10 m 19 s
1000	1 h 22 m 59 s

Acknowledgments

The last author would like to acknowledge partial support from NSF grants SES-0835550 (CDI), CMMI-0844795 (CAREER), CMMI-0968842, a Peek junior faculty fellowship, a Howard B. Wentz Jr. award and a gift from Google.

References

- E. Candes and T. Tao. The Dantzig selector: statistical estimation when p is much larger than n . *Annals of Statistics*, 35(6):2313–2351, 2007.
- A. d’Aspremont, L. El Ghaoui, M.I. Jordan, and G. R. G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *SIAM Review*, 49(3):434–448, 2007.
- A. d’Aspremont, F. Bach, and L. El Ghaoui. Optimal solutions for sparse principal component analysis. *Journal of Machine Learning Research*, 9:1269–1294, 2008.
- D.L. Donoho and J. Tanner. Neighborliness of randomly projected simplices in high dimensions. *Proceedings of the National Academy of Sciences*, 102(27):9452–9457, 2005.
- B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32(2):407–499, 2004.
- M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image Processing*, 15(12):3736–3745, 2006.
- U. Feige and M. Langberg. Approximation algorithms for maximization problems arising in graph partitioning. *Journal of Algorithms*, 41(2):174–211, 2001.
- U. Feige and M. Seltser. On the densest k -subgraph problem. Technical report, Department of Applied Mathematics and Computer Science, The Weizmann Institute, 1997.
- U. Feige, D. Peleg, and G. Kortsarz. The dense k -subgraph problem. *Algorithmica*, 29(3):410–421, 2001.
- G.M. Furnival and R.W. Wilson Jr. Regressions by leaps and bounds. *Technometrics*, 42(1):69–79, 2000.
- M.X. Goemans and D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42:1115–1145, 1995.
- DJ Hand. Branch and bound in statistical data analysis. *The Statistician*, pages 1–13, 1981.
- R.A. Horn and C.R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

- G. Kortsarz and D. Peleg. On choosing a dense subgraph. In *Foundations of Computer Science, 1993. Proceedings., 34th Annual Symposium on*, pages 692–701, 1993.
- J. Mairal, G. Sapiro, and M. Elad. Learning multiscale sparse representations for image and video restoration. *SIAM Multiscale Modeling and Simulation*, 7(1):214–241, 2008.
- N. Meinshausen, G. Rocha, and B. Yu. A tale of three cousins: Lasso, l2boosting, and danzig. *Annals of Statistics*, 35(6):2373–2384, 2007.
- B. Moghaddam, A. Gruber, Y. Weiss, and S. Avidan. Sparse regression as a sparse eigenvalue problem. In *Information Theory and Applications Workshop, 2008*, pages 121–127, 2008.
- PM Narendra and K. Fukunaga. A branch and bound algorithm for feature subset selection. *IEEE Transactions on Computers*, 100(26):917–922, 1977.
- B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, 1995.
- Y. Nesterov. *Introductory Lectures on Convex Optimization*. Springer, 2003.
- R. Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal statistical society, series B*, 58(1):267–288, 1996.