EE364b Homework 6

1. *Conjugate gradient residuals.* Let $r^{(k)} = b - Ax^{(k)}$ be the residual associated with the kth element of the Krylov sequence. Show that $r^{(j)T}r^{(k)} = 0$ for $j \neq k$. In other words, the Krylov sequence residuals are mutually orthogonal. Do not use the explicit algorithm to show this property; use the basic definition of the Krylov sequence, *i.e.*, $x^{(k)}$ minimizes $(1/2)x^T Ax - b^T x$ over $\mathcal{K}_k$.

2. *CG and PCG example.* In this problem you explore a variety of methods to solve $Ax = b$, where $A \in \mathbf{S}^n_{++}$ has block diagonal plus sparse structure: $A = A_{\mathrm{blk}} + A_{\mathrm{sp}}$, where $A_{\mathrm{blk}} \in \mathbf{S}^n_{++}$ is block diagonal and $A_{\mathrm{sp}} \in \mathbf{S}^n$ is sparse. For simplicity we assume A_{blk} consists of k blocks of size m, so $n = mk$. The matrix A_{sp} has N nonzero elements.

   (a) What is the approximate flop count for solving $Ax = b$ if we treat A as dense?

   (b) What is the approximate flop count for an iteration of CG? (Assume multiplication by A_{blk} and A_{sp} are done exploiting their respective structures.) You can ignore the handful of inner products that need to be computed.

   (c) Now suppose that we use PCG, with preconditioner $M = A_{\mathrm{blk}}^{-1}$. What is the approximate flop count for computing the Cholesky factorization of A_{blk}? What is the approximate flop count per iteration of PCG, once a Cholesky factorization of A_{blk} if found?

   (d) Now consider the specific problem with A_{blk}, A, and b generated by `ex_blockprecond.m`. Solve the problem using direct methods, treating A as dense, and also treating A as sparse. Run CG on the problem for a hundred iterations or so, and plot the relative residual versus iteration number. Run PCG on the same problem, using the block-diagonal preconditioner $M = A_{\mathrm{blk}}^{-1}$. Give the solution times for dense direct, sparse direct, CG (to relative residual 10^{-4}, say), and PCG (to relative residual 10^{-8}, say). For PCG break out the time as time for initial Cholesky factorization, and time for PCG iterations.

   Hints.

   - To force Matlab to treat A as dense, use `full(A)`.

   - You do not need to implement the conjugate gradient algorithm; instead use the `pcg` function in Matlab.

   - To block precondition with $M = A_{\mathrm{blk}}^{-1}$, first find the Cholesky factorization of A_{blk}, *i.e.*, lower triangular L with $LL^T = A_{\mathrm{blk}}$. The Matlab code to implement block preconditioning is

   ```
   L = chol(A_blk)';
   [...] = pcg(A,b,tolerance,MAXITER,L,L');
   ```

Matlab uses a sparse Cholesky to factor A_{blk}, which is less efficient than using a dense Cholesky factorization on each block separately, but it is efficient enough to make the point.

3. *Sum-of-norms heuristic for block sparsity.* The basic ℓ_1 heuristic for finding a sparse vector x in a convex set $\mathcal{C}$ is to minimize $\|x\|_1$ over $\mathcal{C}$. We are given a partition of x into subvectors: $x = (x^{(1)}, \ldots, x^{(k)})$, with $x^{(i)} \in \mathbf{R}^{n_i}$. Our goal is to find an $x \in \mathcal{C}$ with the fewest number of subvectors $x^{(i)}$ nonzero. Like the basic sparsity problem, this problem is in general difficult to solve. But a variation on the ℓ_1 heuristic can work well to give a block sparse, if not the block sparsest, vector. The heuristic is to minimize

$$\sum_{i=1}^{k} \|x^{(i)}\|$$

over $x \in \mathcal{C}$.

(a) What happens if the norms above are ℓ_1 norms? Would you expect x to be sparse? Block sparse? (Your answer can be very brief.)

(b) Generate a nonempty polyhedron $\{y \mid Ay \preceq b\}$, with $A \in \mathbf{R}^{50 \times 100}$, and $b \in \mathbf{R}^{50}$. We will divide x into 20 subvectors, each of length 5. Use the heuristic above, with ℓ_1, ℓ_2, and ℓ_∞ norms on the subvectors, to find block sparse x in your polyhedron.

Hints. You may find the functions `norms` and `reshape` useful.

4. *Sparse Bayes network identification.* Suppose you are given samples $y_1, \ldots, y_N \in \mathbf{R}^n$ from an $\mathcal{N}(0, \Sigma)$ distribution, where $\Sigma \succ 0$. (See page 355 of the textbook.) From these samples you need to estimate the parameter Σ. You'd like Σ^{-1} to be sparse, if possible. A sparse Σ^{-1} corresponds to a compact Bayes network representation, since $(\Sigma^{-1})_{ij} = 0$ means that x_i and x_j are conditionally independent, given the other x_k's.

(a) Suggest a simple method for finding Σ with Σ^{-1} sparse, with the likelihood of Σ nearly maximum. Your method should involve a parameter $\lambda \geq 0$ that adjusts the trade-off between suboptimality in likelihood, and sparsity of Σ^{-1}.

(b) Carry out your method on the data found in `sp_bayesnet_data.m`. The true inverse covariance, Σ^{-1}, is stored in the matrix `Siginv`, and the samples $y_1, \ldots, y_N$ are stored in the matrix `Ys`. (You may need to experiment with different values of λ.) Plot the sparsity patterns of Σ^{-1} and your reconstruction $\hat{\Sigma}^{-1}$.

5. *Rank minimization via the nuclear norm.* A simple heuristic for finding a matrix of low (if not minimum) rank, in a convex set, is to minimize its nuclear norm, *i.e.*, the sum of its singular values, over the convex set. Test this method on a numerical example, obtained by generating data matrices $A_0, \ldots, A_n \in \mathbf{R}^{p \times q}$ and attempting to minimize the rank of $A_0 + x_1 A_1 + \cdots + x_n A_n$. You can use $n = 50$, $p = 10$, $q = 10$.

6. *Portfolio investment with linear and fixed costs.* The goal is to optimally invest an initial amount B in a portfolio of n assests. Let x_i be the amount (in dollars) of asset i that we purchase. We assume that no short selling is allowed, *i.e.*, $x \succeq 0$. We pay a fixed cost β for each asset we invest in, *i.e.*, for each asset with $x_i > 0$. We also pay a fee that is proportional to the amount purchased, given by $\alpha^T x$, where α_i is the fee (rate) associated with asset i. The budget constraint can be expressed as

$$\mathbf{1}^T x + \beta \mathbf{card}(x) + \alpha^T x \leq B.$$

The mean return on the portfolio is given by $\mu^T x$, where μ_i is the mean return of asset i, and the variance of the portfolio is $x^T \Sigma x$, where Σ is the covariance of the price changes. Our goal is to minimize the portfolio standard deviation, subject to meeting a minimum mean return requirement. This can be expressed as the problem

$$
\begin{array}{ll}
\text{minimize} & (x^T \Sigma x)^{1/2} \\
\text{subject to} & \mu^T x \geq R_{\min}, \qquad x \succeq 0 \\
& \mathbf{1}^T x + \beta \mathbf{card}(x) + \alpha^T x \leq B.
\end{array}
\tag{1}
$$

This is a convex-cardinality problem. Once the choice is made of which assets to invest in, the problem becomes convex. But there are 2^n possible subsets of assets to invest in, so exhaustive search over all of these is not practical for $n \geq 20$.

You will focus on a particular instance of this problem, with data given in the file `l1_heuristic_portfolio_data.m` on the class web site.

(a) *Heuristic portfolio investment.* You will use the following ℓ_1-norm based heuristic to approximately solve the portfolio investment problem. First we replace $\mathbf{card}(x)$ with $\gamma \mathbf{1}^T x$ (which is the same as $\gamma \|x\|_1$, since $x \succeq 0$), where $\gamma \geq 0$ is a parameter, to get the problem

$$
\begin{array}{ll}
\text{minimize} & (x^T \Sigma x)^{1/2} \\
\text{subject to} & \mu^T x \geq R_{\min}, \qquad x \succeq 0 \\
& \mathbf{1}^T x + \beta \gamma \mathbf{1}^T x + \alpha^T x \leq B.
\end{array}
\tag{2}
$$

Solve (2) for, say, 50 values of γ in $[0, 25]$. Note that when you solve the problem for $\gamma = 0$, you are just ignoring the fixed investment costs. The solutions of (2) need not be feasible for the portfolio investment problem (1), since the true budget constraint can be violated.

For each solution of (2), note the sparsity pattern of x. Fix this sparsity pattern (which makes $\mathbf{card}(x)$ constant), and solve the problem (1). This procedure is called *polishing*.

Plot the portfolio standard deviation obtained (after polishing) versus γ. What is the minimum standard deviation, $\sigma_{\min}$, that you obtain? Give the best portfolio found (*i.e.*, the assets purchased, and the amounts for each one). Also plot $\mathbf{card}(x)$, *i.e.*, the number of assets invested in, versus γ.

3

Hint. To determine the sparsity pattern of x after solving (2), you'll need to use a reasonable (positive) threshold to determine if $x_i = 0$, as in `find(x<1e-3)`.

(b) *A lower bound.* For some values of γ the optimal value of (2) is a lower bound on the optimal value of the original portfolio investment problem (1). Find a simple value $\tilde{\gamma}$ of γ for which this is the case. Compute the corresponding lower bound and compare it to the standard deviation found in (a).

(c) *A more sophisticated lower bound.* Let $x^\star$ denote an optimal point for the original portfolio investment problem (1). Suppose $x_i^\star \le u_i$ for $i = 1, \ldots, n$. Explain why the optimal value of the problem

$$\begin{array}{ll}
\text{minimize} & (x^T \Sigma x)^{1/2} \\
\text{subject to} & \mu^T x \ge R_{\min}, \qquad x \succeq 0 \\
& \mathbf{1}^T x + \beta v^T x + \alpha^T x \le B,
\end{array}$$

where $v_i = 1/u_i$, gives a lower bound on the optimal value of the original portfolio investment problem (1).

A simple choice for u_i is $u_i = B$. Better (*i.e.*, smaller) values can be found as the optimal values of the problems

$$\begin{array}{ll}
\text{maximize} & x_i \\
\text{subject to} & \mu^T x \ge R_{\min}, \qquad x \succeq 0 \\
& (x^T \Sigma x)^{1/2} \le \sigma_{\min} \\
& \mathbf{1}^T x + \beta \tilde{\gamma} \mathbf{1}^T x + \alpha^T x \le B,
\end{array}$$

where $\sigma_{\min}$ is the standard deviation of the portfolio found in part (a). Justify this.

Carry out this procedure for the given problem instance, and compare the resulting lower bound to the results from parts (a) and (b).