On the behavior of the conjugate-gradient method on ill-conditioned problems

Anders Forsgren

Optimization and Systems Theory
Department of Mathematics
Royal Institute of Technology (KTH)
Stockholm, Sweden

13th Mathematical Programming Meeting
Han-sur-Lesse, Belgium
March 16–17, 2006

Linear equation considered

Fundamental problem in linear algebra: Solve $Ax = b$, where A is an $n \times n$ symmetric positive definite matrix and b is an n-dimensional vector.

From an optimization perspective, consider

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \frac{1}{2} x^T H x + c^T x,$$

where H is an $n \times n$ symmetric positive definite matrix and c is an n-dimensional vector. The unique solution is given by $Hx = -c$.

By identifying $A = H$ and $b = -c$, the problems are equivalent.

How to solve the linear equation

We want to solve $Hx = -c$, where $H = H^T \succ 0$.

Two major ways of solving the linear equations:

- Matrix factorization. Method of choice: Cholesky factorization.
- Iterative method. Method of choice: Conjugate-gradient method.

In this talk, the focus is on the conjugate-gradient method.

We are interested in the ill-conditioned case, when the dimension of H is large and the eigenvalues can be divided into two groups: large and small.

Why consider the conjugate-gradient method?

Why is it relevant to consider the conjugate-gradient method in this presentation?

- Important to be able to solve linear equations approximately in a Newton-based method for smooth optimization.
- Gives insight into the behavior of the generated iterates for certain optimization problems.
- Gives insight into how a particular nonlinear optimization problem with several thousand variables can be “solved” in 25 iterations using a quasi-Newton based method.

Motivation for research

Our motivation for considering the conjugate-gradient method is twofold:

- Intensity modulated radiation therapy.
- Solving linear equations that arise in interior methods.

Outline of talk

- The conjugate gradient method.
- Application I: Intensity modulated radiation therapy.
- Application II: Interior methods.

The conjugate gradient method

Algorithm. *The conjugate gradient algorithm*

$k \leftarrow 0$; $x^{(k)} \leftarrow$ initial point; $g^{(k)} \leftarrow Hx^{(k)} + c$;

while $\|g^{(k)}\| \neq 0$

if $k > 0$ **then**

$$\beta^{(k)} \leftarrow \frac{g^{(k)T} g^{(k)}}{g^{(k-1)T} g^{(k-1)}}; \quad p^{(k)} \leftarrow -g^{(k)} + \beta^{(k)} p^{(k-1)};$$

else

$$p^{(k)} \leftarrow -g^{(k)};$$

end if

$$\alpha^{(k)} \leftarrow -\frac{g^{(k)T} p^{(k)}}{p^{(k)T} H p^{(k)}};$$

$$x^{(k+1)} \leftarrow x^{(k)} + \alpha^{(k)} p^{(k)}; \quad g^{(k+1)} \leftarrow g^{(k)} + \alpha^{(k)} H p^{(k)};$$

$$k \leftarrow k + 1;$$

end

A reformulation of the quadratic problem

The aim is to solve

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \tfrac{1}{2} x^T H x + c^T x.$$

Let $x^* = -H^{-1} c$ denote the optimal solution.

Then, $\frac{1}{2} x^T H x + c^T x = \frac{1}{2}(x - x^*)^T H (x - x^*) - \frac{1}{2} x^{*T} H x^*$.

With $\xi = x - x^*$, the problem may equivalently be rewritten as

$$\underset{\xi \in \mathbb{R}^n}{\text{minimize}} \quad \tfrac{1}{2} \xi^T H \xi.$$

We refer to ξ as the *residual*.

Note that x^* is not known, but we may still discuss properties of ξ.

The first conjugate-gradient iteration

The first conjugate-gradient iteration is a steepest-descent iteration.

Given $\xi^{(0)}$, we compute $g^{(0)} = H\xi^{(0)}$ and solve

$$\begin{array}{ll} \underset{\xi \in \mathbb{R}^n, \eta \in \mathbb{R}}{\text{minimize}} & \frac{1}{2}\xi^T H \xi \\ \text{subject to} & \xi = \xi^{(0)} - g^{(0)}\eta. \end{array}$$

Optimality conditions given by

$$\begin{aligned} H\xi - g &= 0, \\ g^{(0)T} g &= 0, \\ \xi &= \xi^{(0)} - g^{(0)}\eta. \end{aligned}$$

Solution $\xi^{(1)}$, $g^{(1)}$ and $\eta^{(1)}$.

The second conjugate-gradient iteration

If $g^{(1)} \neq 0$, we make a two-dimensional search and solve

$$\begin{array}{ll} \underset{\xi \in \mathbb{R}^n, \eta \in \mathbb{R}^2}{\text{minimize}} & \frac{1}{2}\xi^T H \xi \\ \text{subject to} & \xi = \xi^{(0)} - g^{(0)}\eta_1 - g^{(1)}\eta_2. \end{array}$$

Optimality conditions given by

$$\begin{aligned} H\xi - g &= 0, \\ g^{(0)T} g &= 0, \\ g^{(1)T} g &= 0, \\ \xi &= \xi^{(0)} - g^{(0)}\eta_1 - g^{(1)}\eta_2. \end{aligned}$$

Solution $\xi^{(2)}$, $g^{(2)}$ and $\eta^{(2)}$.

The kth conjugate-gradient iteration

If $g^{(k-1)} \neq 0$, we make a k-dimensional search and solve

$$\begin{array}{ll} \underset{\xi\in I\!\!R^n,\eta\in I\!\!R^k}{\text{minimize}} & \frac{1}{2}\xi^T H\xi \\ \text{subject to} & \xi = \xi^{(0)} - \sum_{l=1}^k g^{(l-1)}\eta_l. \end{array}$$

Optimality conditions given by

$$\begin{aligned} H\xi - g &= 0, \\ g^{(l)T} g &= 0, \quad l = 0, \dots, k-1, \\ \xi &= \xi^{(0)} - \sum_{l=1}^{k} g^{(l-1)}\eta_l. \end{aligned}$$

Solution $\xi^{(k)}$, $g^{(k)}$ and $\eta^{(k)}$.

Recurrence relation given by the conjugate-gradient algorithm.

A Krylov vector representation

Note that

$$g^{(0)} = H\xi^{(0)},$$
$$g^{(1)} = H\xi^{(1)} = H\xi^{(0)} - H^2\xi^{(0)}\eta^{(1)}, \quad \text{etc.}$$

We may replace $\xi = \xi^{(0)} - \sum_{l=1}^{k} g^{(l-1)}\eta_l$ by $\xi = \xi^{(0)} + \sum_{l=1}^{k} H^l\xi^{(0)}\gamma_l$.

The conjugate-gradient subproblem may then be written as

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\gamma\in I\!R^k}{\text{minimize}} & \frac{1}{2}\xi^T H\xi \\ \text{subject to} & \xi = \xi^{(0)} + \sum_{l=1}^{k} H^l\xi^{(0)}\gamma_l. \end{array}$$

This formulation is a convex quadratic program, where the *Krylov vectors* $\xi^{(0)}, H\xi^{(0)}, \ldots, H^k\xi^{(0)}$ appear explicitly.

A quadratic-programming conjugate gradient subproblem

In the problem

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\gamma\in I\!R^k}{\text{minimize}} & \frac{1}{2}\xi^T H \xi \\ \text{subject to} & \xi = \xi^{(0)} + \sum_{l=1}^{k} H^l \xi^{(0)} \gamma_l, \end{array}$$

we may consider a rotation of the variables so that the Hessian becomes diagonal, i.e., $H = \text{diag}(\lambda)$, where λ is the vector of eigenvalues of H. We order the eigenvalues such that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n > 0$.

The conjugate-gradient method then gives $\xi^{(k)}$ and $\gamma^{(k)}$ from

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\gamma\in I\!R^k}{\text{minimize}} & \frac{1}{2}\displaystyle\sum_{i=1}^{n} \lambda_i \xi_i^2 \\ \text{subject to} & \xi_i = \xi_i^{(0)} + \sum_{l=1}^{k} \lambda_i^l \xi_i^{(0)} \gamma_l, \quad i = 1, \dots, n. \end{array}$$

A polynomial conjugate gradient subproblem

Alternatively, note that for $i = 1, \ldots, n$,

$$\xi_i = \xi_i^{(0)} + \sum_{l=1}^{k} \lambda_i^l \xi_i^{(0)} \gamma_l = \left(1 + \sum_{l=1}^{k} \lambda_i^l \gamma_l \right) \xi_i^{(0)},$$

may be represented as $\xi_i = P_k(\lambda_i)\xi^{(0)}$, where $P_k(\lambda)$ is a kth degree polynomial with $P_k(0) = 1$.

We may characterize $P_k(\lambda)$ in terms of its zeros $\zeta \in I\!R^k$ as

$$Q_k(\lambda, \zeta) = \prod_{l=1}^{k} \left(1 - \frac{\lambda}{\zeta_l} \right),$$

where we choose the ordering $\zeta_1 \geq \zeta_2 \geq \cdots \geq \zeta_k$.

A polynomial conjugate gradient subproblem

The conjugate gradient subproblem may then equivalently be written as

$$\begin{array}{ll} \underset{\xi\in\mathbb{R}^n,\zeta\in\mathbb{R}^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{n}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i = Q_k(\lambda_i,\zeta)\xi_i^{(0)}, \quad i=1,\dots,n, \end{array}$$

where the optimal solution is denoted by $\xi^{(k)}$ and $\zeta^{(k)}$.

Equivalently,

$$\begin{array}{ll} \underset{\xi\in\mathbb{R}^n,\zeta\in\mathbb{R}^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{n}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i = \prod_{l=1}^{k}\left(1-\frac{\lambda_i}{\zeta_l}\right)\xi_i^{(0)}, \quad i=1,\dots,n. \end{array}$$

Conjugate gradient subproblems

Two equivalent characterizations of the conjugate gradient subproblem:

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\gamma\in I\!R^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{n}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i = \xi_i^{(0)} + \sum_{l=1}^{k}\lambda_i^l\xi_i^{(0)}\gamma_l, \quad i=1,\dots,n. \end{array}$$

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\zeta\in I\!R^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{n}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i = \prod_{l=1}^{k}\left(1-\frac{\lambda_i}{\zeta_l}\right)\xi_i^{(0)}, \quad i=1,\dots,n. \end{array}$$

Well-conditioned example problem

Example problem with $\lambda = (2, 1.5, 1)^T$ and $\xi^{(0)} = (1, 1, 1)^T$.

	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$
$i = 1$	2.0000	1.0000	-0.1717	0.0180	0.0000
$i = 2$	1.5000	1.0000	0.1212	-0.0641	0.0000
$i = 3$	1.0000	1.0000	0.4141	0.0721	-0.0000

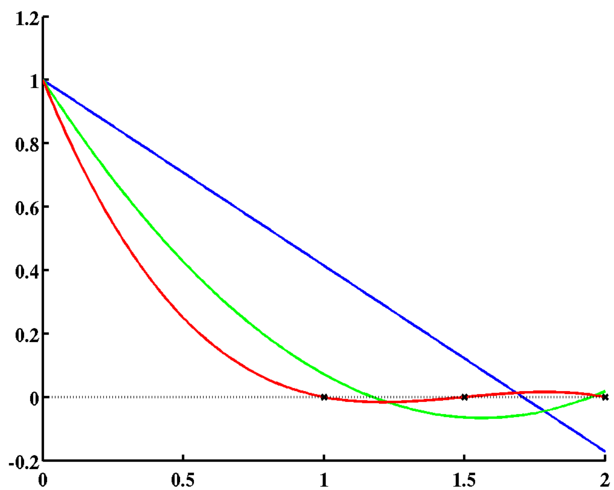

Polynomials for well-conditioned example problem

Example problem with $\lambda = (2, 1.5, 1)^T$ and $\xi^{(0)} = (1, 1, 1)^T$.

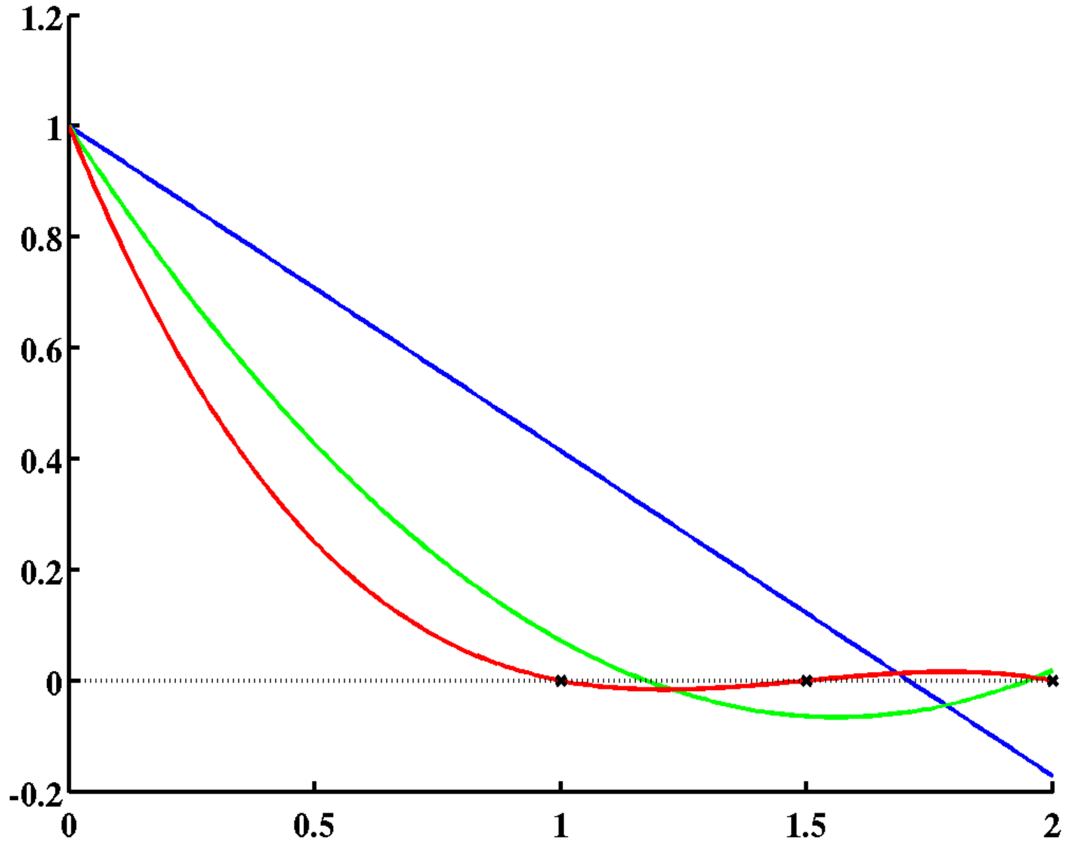

Behavior of the conjugate gradient subproblems

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\zeta\in I\!R^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{n}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i = \prod_{l=1}^{k}\left(1-\frac{\lambda_i}{\zeta_l}\right)\xi_i^{(0)}, \quad i=1,\dots,n, \end{array}$$

The optimal solution $\xi^{(k)}$ will tend to have smaller components $\xi_i^{(k)}$ for i such that λ_i is large and/or $\xi_i^{(0)}$ is large.

Nonlinear dependency of $\xi^{(k)}$ on λ and $\xi^{(0)}$.

We are interested in the ill-conditioned case, when the dimension of H is large and the eigenvalues can be divided into two groups: large and small.

Residuals for ill-conditioned example problem I

Problem with $\lambda = (2, 1.5, 1, 0.1, 0.01)^T$ and $\xi^{(0)} = (1, 1, 1, 1, 1)^T$.

i	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$	$\xi^{(4)}$	$\xi^{(5)}$
1	2.0000	1.0000	-0.1733	0.0217	-0.0049	0.0000	0.0000
2	1.5000	1.0000	0.1201	-0.0702	0.0236	-0.0000	-0.0000
3	1.0000	1.0000	0.4134	0.0623	-0.0413	0.0001	-0.0000
4	0.1000	1.0000	0.9413	0.8659	0.7647	-0.0108	0.0000
5	0.0100	1.0000	0.9941	0.9862	0.9747	0.8793	-0.0000

	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$
$i = 1$	2.0000	1.0000	-0.1717	0.0180	0.0000
$i = 2$	1.5000	1.0000	0.1212	-0.0641	0.0000
$i = 3$	1.0000	1.0000	0.4141	0.0721	-0.0000

Empirical behavior on ill-conditioned problem

Based on the ill-conditioned example, we make the following observations:

- The conjugate-gradient method initially "approximately" solves the problem given by the large eigenvalues only.
- During this initial stage, the method "almost" behaves as the conjugate method applied to the smaller problem given by the large eigenvalues only.
- During this initial stage, the residuals associated with small eigenvalues of H are "almost" unchanged.
- The ill-conditioning of the problem does not appear in this initial stage.

Polynomials for ill-conditioned example problem I

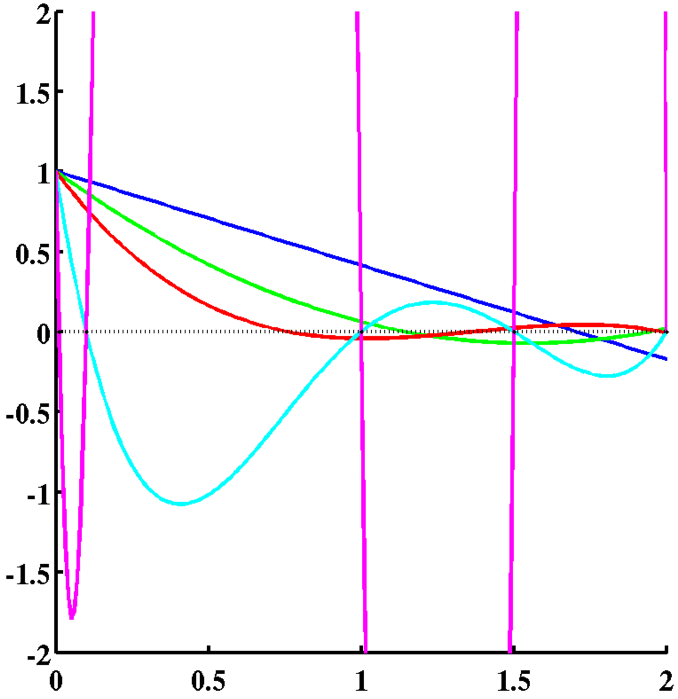

Polynomials $Q^{(k)}(\lambda, \zeta^{(k)})$, $k = 1, \dots, 5$, as a function of λ, for problem with $\lambda = (2, 1.5, 1, 0.1, 0.01)^T$ and $\xi^{(0)} = (1, 1, 1, 1, 1)^T$.

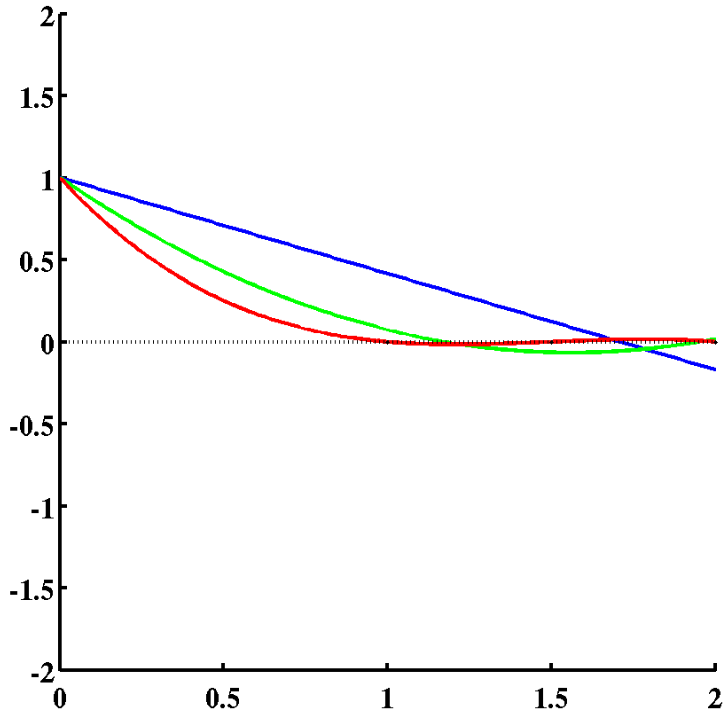

Polynomials $Q^{(k)}(\lambda, \zeta^{(k)})$, $k = 1, \dots, 3$, as a function of λ, for problem with $\lambda = (2, 1.5, 1)^T$ and $\xi^{(0)} = (1, 1, 1)^T$.

Residuals for ill-conditioned example problem II

Problem with $\lambda = (2, 1.5, 1, 0.01, 0.0001)^T$ and $\xi^{(0)} = (1, 1, 1, 1, 1)^T$.

i	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$	$\xi^{(4)}$	$\xi^{(5)}$
1	2.0000	1.0000	-0.1717	0.0181	-0.0001	0.0000	0.0000
2	1.5000	1.0000	0.1212	-0.0642	0.0003	-0.0000	0.0000
3	1.0000	1.0000	0.4141	0.0720	-0.0006	0.0000	0.0000
4	0.0100	1.0000	0.9941	0.9864	0.9784	-0.0001	0.0000
5	0.0001	1.0000	0.9999	0.9999	0.9998	0.9898	-0.0000

	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$
$i = 1$	2.0000	1.0000	-0.1717	0.0180	0.0000
$i = 2$	1.5000	1.0000	0.1212	-0.0641	0.0000
$i = 3$	1.0000	1.0000	0.4141	0.0721	-0.0000

Residuals for ill-conditioned example problem III

Problem with $\lambda = (2, 1.5, 1, 0.1, 0.01)^T$ and $\xi^{(0)} = (1, 1, 1, 10, 10)^T$.

i	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$	$\xi^{(4)}$	$\xi^{(5)}$
1	2.0000	1.0000	-0.3242	0.3162	-0.0964	0.0003	0.0000
2	1.5000	1.0000	0.0068	-0.5540	0.4646	-0.0020	0.0000
3	1.0000	1.0000	0.3379	-0.7300	-0.8113	0.0052	-0.0000
4	0.1000	10.0000	9.3379	7.0207	1.4241	-0.1078	-0.0000
5	0.0100	10.0000	9.9338	9.6896	9.0454	8.7927	0.0000

	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$
$i = 1$	2.0000	1.0000	-0.1717	0.0180	0.0000
$i = 2$	1.5000	1.0000	0.1212	-0.0641	0.0000
$i = 3$	1.0000	1.0000	0.4141	0.0721	-0.0000

Residuals for ill-conditioned example problem IV

Problem with $\lambda = (2, 1.5, 1, 0.01, 0.0001)^T$ and $\xi^{(0)} = (1, 1, 1, 10, 10)^T$.

i	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$	$\xi^{(4)}$	$\xi^{(5)}$
1	2.0000	1.0000	-0.1733	0.0225	-0.0072	0.0000	-0.0000
2	1.5000	1.0000	0.1200	-0.0712	0.0341	-0.0000	-0.0000
3	1.0000	1.0000	0.4133	0.0605	-0.0577	0.0000	0.0000
4	0.0100	10.0000	9.9413	9.8614	9.7316	-0.0010	-0.0000
5	0.0001	10.0000	9.9994	9.9986	9.9973	9.8978	0.0000

	λ	$\xi^{(0)}$	$\xi^{(1)}$	$\xi^{(2)}$	$\xi^{(3)}$
$i = 1$	2.0000	1.0000	-0.1717	0.0180	0.0000
$i = 2$	1.5000	1.0000	0.1212	-0.0641	0.0000
$i = 3$	1.0000	1.0000	0.4141	0.0721	-0.0000

Quantification of empirical results

We consider the case when H has r large eigenvalues and the remaining $n - r$ eigenvalues are small.

Let $\epsilon \in (0, 1]$ be such that $\lambda_{r+1} \leq \epsilon\lambda_r$. We will be interested in the case when $\epsilon \ll 1$. The analysis applies when $\epsilon r \leq 1$.

The basis for our analysis is to ignore eigenvalues $r + 1$ through n and consider the problem

$$\begin{array}{ll} \underset{\xi \in I\!R^r, \gamma \in I\!R^k}{\text{minimize}} & \frac{1}{2} \sum_{i=1}^{r} \lambda_i \xi_i^2 \\ \text{subject to} & \xi_i = \xi_i^{(0)} + \sum_{l=1}^{k} \lambda_i^l \xi_i^{(0)} \gamma_l, \quad i = 1, \ldots, r, \end{array}$$

where we denote the optimal solution by $\bar{\xi}_i^{(k)}$, $i = 1, \ldots, r$, and $\bar{\gamma}^{(k)}$.

A relaxed problem

Computing $\bar{\xi}_i^{(k)}$, $i = 1, \dots, r$, and $\bar{\gamma}^{(k)}$ from

$$\begin{array}{ll} \underset{\xi \in I\!R^r, \gamma \in I\!R^k}{\text{minimize}} & \frac{1}{2} \sum_{i=1}^{r} \lambda_i \xi_i^2 \\ \text{subject to} & \xi_i = \xi_i^{(0)} + \sum_{l=1}^{k} \lambda_i^l \xi_i^{(0)} \gamma_l, \quad i = 1, \dots, r, \end{array}$$

and setting $\bar{\xi}_i^{(k)} = \xi_i^{(0)} + \sum_{l=1}^{k} \lambda_i^l \xi_i^{(0)} \bar{\gamma}_l^{(k)}$, $i = r+1, \dots, n$, is equivalent to solving

$$\begin{array}{ll} \underset{\xi \in I\!R^n, \gamma \in I\!R^k}{\text{minimize}} & \frac{1}{2} \sum_{i=1}^{r} \lambda_i \xi_i^2 \\ \text{subject to} & \xi_i = \xi_i^{(0)} + \sum_{l=1}^{k} \lambda_i^l \xi_i^{(0)} \gamma_l, \quad i = 1, \dots, n. \end{array}$$

We refer to this problem as the *relaxed problem*.

The original problem and the relaxed problem

The original problem may be written as

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\gamma\in I\!R^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{r}\lambda_i\xi_i^2+\frac{1}{2}\sum_{i=r+1}^{n}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i=\xi_i^{(0)}+\sum_{l=1}^{k}\lambda_i^l\xi_i^{(0)}\gamma_l, \quad i=1,\dots,n. \end{array}$$

and the relaxed problem may be written as

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\gamma\in I\!R^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{r}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i=\xi_i^{(0)}+\sum_{l=1}^{k}\lambda_i^l\xi_i^{(0)}\gamma_l, \quad i=1,\dots,n. \end{array}$$

The term $\frac{1}{2}\sum_{i=r+1}^{n}\lambda_i\xi_i^2$ is omitted in the relaxed problem.

Equivalent relaxed problems

The relaxed problem may be formulated either as

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\gamma\in I\!R^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{r}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i = \xi_i^{(0)} + \sum_{l=1}^{k}\lambda_i^l\xi_i^{(0)}\gamma_l, \quad i=1,\dots,n, \end{array}$$

or as

$$\begin{array}{ll} \underset{\xi\in I\!R^n,\zeta\in I\!R^k}{\text{minimize}} & \frac{1}{2}\sum_{i=1}^{r}\lambda_i\xi_i^2 \\ \text{subject to} & \xi_i = \prod_{l=1}^{k}\left(1-\frac{\lambda_i}{\zeta_l}\right)\xi_i^{(0)}, \quad i=1,\dots,n, \end{array}$$

where the optimal solutions are denoted by $\bar{\xi}^{(k)}$, $\bar{\gamma}^{(k)}$ and $\bar{\zeta}^{(k)}$.

Note that the problem only concerns the first r eigenvalues.

We refer to iterations 1 through r as *early iterations*.

Behavior on the relaxed problem for early iterations

We initially show that for the relaxed problem, the residuals associated with the small eigenvalues are not altered much in the early iterations.

Lemma. *Assume that* $\lambda_{r+1} \le \epsilon\lambda_r$. *Then, for* $k = 1, \ldots, r$ *and* $i = r+1, \ldots, n$,

$$\begin{aligned} \xi_i^{(0)} \ge \bar{\xi}_i^{(k)} \ge (1-k\epsilon)\xi_i^{(0)}, &\quad \textit{if } \xi_i^{(0)} \ge 0, \\ \xi_i^{(0)} \le \bar{\xi}_i^{(k)} \le (1-k\epsilon)\xi_i^{(0)}, &\quad \textit{if } \xi_i^{(0)} \le 0. \end{aligned}$$

This is a consequence of $Q_k(\lambda, \bar{\zeta}^{(k)})$ being convex for $\lambda \in (0, \bar{\zeta}_k^{(k)})$, in cunjunction with $\bar{\zeta}^{(k)} \in [\lambda_k, \lambda_1]$ for $k = 1, \ldots, k$.

Behavior of the residuals associated with large eigenvalues

We can now characterize the behavior of the residuals associated with the large eigenvalues.

Theorem. *Assume that* $\lambda_{r+1} \le \epsilon\lambda_r$ *and* $\epsilon r \le 1$. *Then,*

$$\sum_{i=1}^{r}(\xi_i^{(k)} - \bar{\xi}_i^{(k)})^2 \le \epsilon \sum_{i=r+1}^{n}(\xi_i^{(0)})^2, \quad k = 1, \ldots, r,$$

$$\sum_{i=1}^{r}(\xi_i^{(k)})^2 \le \epsilon \sum_{i=r+1}^{n}(\xi_i^{(0)})^2, \quad k = r+1, \ldots, n.$$

The bound is obtained via the relaxed problem.

Behavior of the residuals associated with small eigenvalues

We can also characterize the behavior of the residuals associated with the small eigenvalues for the early iterations.

Theorem. *Assume that* $\lambda_{r+1} \le \epsilon\lambda_r$ *and* $\epsilon r \le 1$. *Let* $\Xi_L^{(0)} = \text{diag}(\xi_1^{(0)}, \dots, \xi_r^{(0)})$, *let* $\xi_S^{(0)} = (\xi_{r+1}^{(0)}, \dots, \xi_n^{(0)})^T$, *and let* $V_L^{(k)}$ *be the* $r \times k$ *matrix with element* ij *given by* $(V_L^{(k)})_{ij} = (\lambda_i/\lambda_1)^j$. *Then, for* $k = 1, \dots, r$,

$$|\xi_i^{(k)} - \bar{\xi}_i^{(k)}| \le \frac{k^{1/2}\epsilon^{3/2}\|\xi_S^{(0)}\|}{\sigma_k(\Xi_L^{(0)} V_L^{(k)})} |\xi_i^{(0)}|, \quad i = r+1, \dots, n,$$

where σ_k *denotes the* k*th singular value.*

This bound is not as explicit.

Summary of the behavior of the conjugate gradient iterates

We have demonstrated how the conjugate gradient method on an ill-conditioned problem generates iterates that are close to the iterates that are generated if only the large eigenvalues are considered.

In particular, the ill-conditioning of the problem does not appear in the early iterations.

This is a desirable feature in some applications. Referred to as *iterative regularization*. The regularization parameter is the number of conjugate-gradient iterations to perform.

In other applications, this behavior is not desirable. Preconditioning is typically needed.

Relationship to the steepest-descent method

As a remark, we also briefly review the steepest descent method in a polynomial framework.

The steepest-descent method may be viewed as a conjugate-gradient method which is restarted every iteration.

We obtain $\xi^{(k)}$ and ζ^k as the optimal solution to the problem

$$\begin{array}{ll} \displaystyle\operatorname*{minimize}_{\xi \in \mathbb{R}^n, \zeta \in \mathbb{R}} & \frac{1}{2} \displaystyle\sum_{i=1}^{n} \lambda_i \xi_i^2 \\ \text{subject to} & \xi_i = \left(1 - \dfrac{\lambda_i}{\zeta}\right) \xi_i^{(k-1)}, \quad i = 1, \dots, n. \end{array}$$

Relationship to the steepest-descent method, cont.

The steepest-descent problem has the closed-form solution

$$\zeta^{(k)} = \sum_{j=1}^{n} \left(\frac{\lambda_j^2 (\xi_j^{(k-1)})^2}{\sum_{l=1}^{n} \lambda_l^2 (\xi_l^{(k-1)})^2} \right) \lambda_j, \quad \text{and}$$

$$\xi_i^{(k)} = \left(1 - \frac{\lambda_i}{\zeta^{(k)}}\right) \xi_i^{(k-1)} = \prod_{l=1}^{k} \left(1 - \frac{\lambda_i}{\zeta^{(l)}}\right) \xi_i^{(0)}, \quad i = 1, \ldots, n,$$

Steepest descent forms polynomials by adding one zero at the time, and not changing the zeros that have already been obtained.

Relationship to the steepest-descent method, cont.

First ten polynomials generated by the steepest-descent method for problem with $\lambda = (2, 1.5, 1, 0.1, 0.01)^T$ and $\xi^{(0)} = (1, 1, 1, 1, 1)^T$.

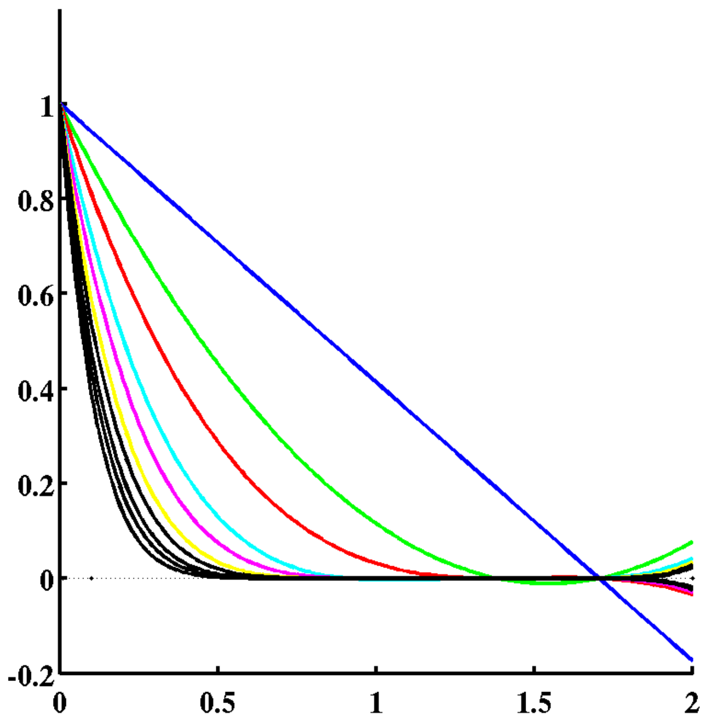

The BFGS quasi-Newton method

The BFGS quasi-Newton method for minimizing the quadratic function takes the following form.

Algorithm. *The BFGS quasi-Newton method*

$k \leftarrow 0$; $x^{(k)} \leftarrow$ initial point; $g^{(k)} \leftarrow Hx^{(k)} + c$; $B^{(0)} \leftarrow I$;

while $\|g^{(k)}\| \neq 0$

Solve $B^{(k)}p^{(k)} = -g^{(k)}$;

$$\alpha^{(k)} \leftarrow -\frac{g^{(k)T}p^{(k)}}{p^{(k)T}Hp^{(k)}};$$

$x^{(k+1)} \leftarrow x^{(k)} + \alpha^{(k)}p^{(k)}$; $g^{(k+1)} \leftarrow g^{(k)} + \alpha^{(k)}Hp^{(k)}$;

$$B^{(k+1)} \leftarrow B^{(k)} - \frac{B^{(k)}p^{(k)}p^{(k)T}B^{(k)}}{p^{(k)T}B^{(k)}p^{(k)}} + \frac{Hp^{(k)}p^{(k)T}H}{p^{(k)T}Hp^{(k)}};$$

$k \leftarrow k + 1$;

end

Methods of conjugate-gradients, quasi-Newton and Lanczos

Consider minimizing a convex quadratic function with exact linsearch.

The conjugate-gradient method and the BFGS quasi-Newton method generate identical search directions $p^{(k)}$.

Quasi-Newton methods and the conjugate-gradient method generate identical iterates $\xi^{(k)}$.

Close connection to Lanczos method for generating eigenvalues. The matrix $g^{(i)T} H g^{(j)}$ is tri-diagonal. The Lanczos vectors are normalized $g^{(j)}$s.

Reference on the behavior of the conjugate-gradient method

Reference:

- A. Forsgren, *On the behavior of the conjugate-gradient method on ill-conditioned problems*. Report TRITA-MAT-2006-OS1, Department of Mathematics, Royal Institute of Technology, 2006.

(Available from http://www.math.kth.se/~andersf)

Application I: Intensity modulated radiation therapy

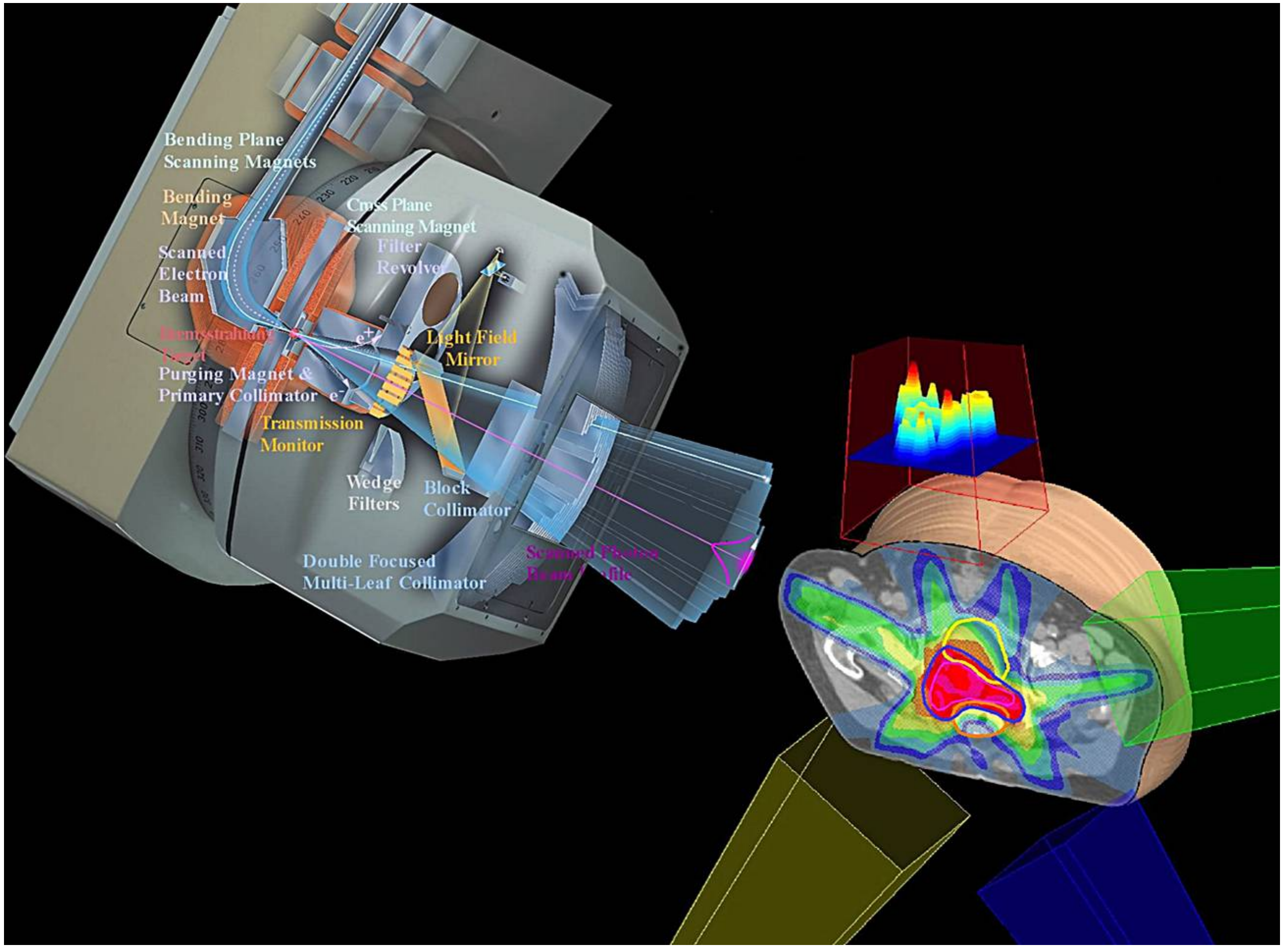

Joint project, KTH and RaySearch Laboratories.

Industrial graduate student: Fredrik Carlsson.

Aim of radiation

The aim of the radiation is typically to give a treatment that leads to a desirable dose distribution in the patient.

Typically, high dose is desired in the tumor cells, and low dose in the other cells.

In particular, certain organs are very sensitive to radiation and must have a low dose level, e.g., the spine.

Hence, a desired dose distribution can be specified, and the question is how to achieve this distribution.

This is an *inverse problem* in that the desired result of the radiation is known, but the treatment plan has to be designed.

Formulation of optimization problem

A radiation treatment is typically given as a series of radiations.

For an individual treatment, the performance depends on

- the beam angle of incidence, which is governed by the supporting gantry; and
- the intensity modulation of the beam, which is governed by the treatment head.

One may now formulate the problem of giving an optimal treatment as an optimization problem, where the variables are the beam angles of incidence and the intensity modulations of the beams.

In this talk, we assume that the gantry angles are fixed on beforehand.

Characteristics of the optimization problem

The resulting optimization problem is a *nonlinear optimization problem*, i.e., it may be posed as the minimization of a smooth nonlinear function subject to smooth constraints.

A feature of the problem is that it is a large-scale problem, typically with a large number of degrees of freedom at the solution.

Many different objective functions have been proposed. The above problem characteristics hold.

Particular optimization problem

Simplified problem takes the form

$$\begin{array}{ll} \underset{x \in \mathbb{R}^n}{\text{minimize}} & \frac{1}{2}\|D(x)Px - \hat{d}\|_2^2 \\ \text{subject to} & x \geq 0. \end{array}$$

where P has dimension of the order 50,000 times 1,000, and $D(x)$ is positive definite and diagonal.

The problem can be solved with a quasi-Newton based sequential-quadratic-programming approach in very few iterations, say 25.

How can that be?

Properties of the optimization problem

The answer lies in the problem characteristics of

$$\underset{x \in I\!R^n, x \geq 0}{\text{minimize}} \quad \frac{1}{2}\|Px - \hat{d}\|_2^2.$$

The matrix P^TP has a low number of dominant, large, eigenvalues. The corresponding eigenvectors yield smooth solutions.

The eigenvectors corresponding to small eigenvalues yield jagged solutions. Undesirable, since they create solutions that are hard to deliver.

Mathematical solution: $x = (P^TP)^{-1}P^T\hat{d}$.

Not the desired solution.

Properties of the optimization problem

A quasi-Newton based sequential-quadratic-programming method for solving

$$\underset{x \in \mathbb{R}^n, x \geq 0}{\text{minimize}} \quad \tfrac{1}{2}\|Px - \hat{d}\|_2^2,$$

will initially tend to proceed in the directions of the eigenvectors corresponding to the large eigenvalues.

The effect is that a high-quality solution with desired properties is obtained in few iterations.

Better, and cheaper to obtain, than mathematically optimal solution.

Quasi-Newton methods equivalent to conjugate gradient method on unconstrained quadratic problem.

Properties of the conjugate-gradient method beneficial in this case.

Reference on intensity modulated radiation therapy

Reference:

- F. Carlsson and A. Forsgren, *Iterative regularization in intensity-modulated radiation therapy optimization*, Medical Physics 33 (2006), 225-234.

(Available from http://www.math.kth.se/~andersf)

Application II: Interior methods for nonlinear optimization

Joint research with Philip E. Gill (UCSD) and Joshua D. Griffin (Sandia).

Solution of linear equations of the form

$$B = \begin{pmatrix} H & A^T \\ A & -G \end{pmatrix},$$

with A an $m \times n$ matrix, H symmetric and G symmetric positive semidefinite.

We will consider the case when G is diagonal and ill-conditioned.

Iterative methods for solving with B interesting since they allow inexact solution to be computed.

Equivalent systems of linear equations

We consider the *augmented system*

$$\begin{pmatrix} H & A^T \\ A & -D \end{pmatrix} \begin{pmatrix} x_1 \\ -x_2 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix},$$

the *condensed system*

$$(H + A^T D^{-1} A) x_1 = b_1 + A^T D^{-1} b_2 \quad \text{and} \quad x_2 = D^{-1}(b_2 - A x_1),$$

or the *doubly augmented system*

$$\begin{pmatrix} H + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} b_1 + 2A^T D^{-1} b_2 \\ b_2 \end{pmatrix},$$

which are all equivalent.

Matrix inertias

The condensed system

$$(H + A^T D^{-1} A) x_1 = b_1 + A^T D^{-1} b_2 \quad \text{and} \quad x_2 = D^{-1}(b_2 - A x_1),$$

and the doubly augmented system

$$\begin{pmatrix} H + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} b_1 + 2A^T D^{-1} b_2 \\ b_2 \end{pmatrix},$$

have matrices that are normally positive definite.

The conjugate gradient method can be applied. *Preconditioning* required.

Projection method required for the augmented system.

The preconditioned conjugate gradient method

The preconditioned conjugate gradient method is formed for a preconditioning matrix M, which is "similar" to H, but for which a solution of $My = h$ is "cheap".

Let $M = R^T R$.

The conjugate gradient method solves $Hx = -c$.

The preconditioned conjugate gradient method solves $R^{-T} H R^{-1} \widetilde{x} = -R^{-T} c$ by the conjugate gradient method. Then, $x = R\widetilde{x}$. The eigenvalues of $R^{-T} H R^{-1}$ determine convergence.

In practice, R is not needed. The eigenvalues of $R^{-T} H R^{-1}$ and $M^{-1} H$ are identical.

Equivalent to quasi-Newton with initial Hessian estimate M.

The preconditioned conjugate gradient method

Algorithm. *The preconditioned conjugate gradient algorithm*

$k \leftarrow 0$; $\quad g^{(k)} \leftarrow Hx^{(k)} + c$; $\quad \widetilde{g}^{(k)} \leftarrow R^{-T} g^{(k)}$;

while $\|\widetilde{g}^{(k)}\| \neq 0$

if $k > 0$ **then**

$$\beta^{(k)} \leftarrow \frac{\widetilde{g}^{(k)T}\widetilde{g}^{(k)}}{\widetilde{g}^{(k-1)T}\widetilde{g}^{(k-1)}}; \quad \widetilde{p}^{(k)} \leftarrow -\widetilde{g}^{(k)} + \beta^{(k)}\widetilde{p}^{(k-1)};$$

else

$$\widetilde{p}^{(k)} \leftarrow -\widetilde{g}^{(k)};$$

end if

$$p^{(k)} \leftarrow R^{-1}\widetilde{p}^{(k)}; \quad \alpha^{(k)} \leftarrow -\frac{\widetilde{g}^{(k)T}\widetilde{p}^{(k)}}{p^{(k)T}Hp^{(k)}};$$

$$\widetilde{x}^{(k+1)} \leftarrow \widetilde{x}^{(k)} + \alpha^{(k)}\widetilde{p}^{(k)}; \quad g^{(k+1)} \leftarrow g^{(k)} + \alpha^{(k)}Hp^{(k)};$$

$$\widetilde{g}^{(k+1)} \leftarrow R^{-T}g^{(k+1)}; \quad k \leftarrow k+1;$$

end

The preconditioned conjugate gradient method

Algorithm. *The preconditioned conjugate gradient algorithm*

$k \leftarrow 0$; $\quad g^{(k)} \leftarrow Hx^{(k)} + c$; $\quad \widetilde{g}^{(k)} \leftarrow M^{-1} g^{(k)}$;

while $\|\widetilde{g}^{(k)}\| \neq 0$

if $k > 0$ **then**

$$\beta^{(k)} \leftarrow \frac{\widetilde{g}^{(k)T} g^{(k)}}{\widetilde{g}^{(k-1)T} g^{(k-1)}}; \quad p^{(k)} \leftarrow -\widetilde{g}^{(k)} + \beta^{(k)} p^{(k-1)};$$

else

$$p^{(k)} \leftarrow -\widetilde{g}^{(k)};$$

end if

$$\alpha^{(k)} \leftarrow -\frac{g^{(k)T} p^{(k)}}{p^{(k)T} H p^{(k)}};$$

$$x^{(k+1)} \leftarrow x^{(k)} + \alpha^{(k)} p^{(k)}; \quad g^{(k+1)} \leftarrow g^{(k)} + \alpha^{(k)} H p^{(k)};$$

$$\widetilde{g}^{(k+1)} \leftarrow M^{-1} g^{(k+1)}; \quad k \leftarrow k + 1;$$

end

Constraint preconditioning

We may use a preconditioning matrix of the form

$$(M + A^T D^{-1} A)$$

for the condensed system and

$$\begin{pmatrix} M + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix}$$

for the doubly augmented system, i.e., H is approximated by M.

Referred to as constraint preconditioning.

Requirement: $M = M^T$, $M + A^T D^{-1} A \succ 0$.

Ideally: M a good approximation to H, and simpler to solve with M.

Solving with the preconditioned matrix

Although the preconditiners are defined for the condensed system

$$(M + A^T D^{-1} A)$$

and the doubly augmented system

$$\begin{pmatrix} M + 2A^T D^{-1} A & A^T \\ A & D \end{pmatrix}$$

we need only solve with the augmented system

$$\begin{pmatrix} M & A^T \\ A & -D \end{pmatrix}.$$

Typically, this can be done via a symmetric indefinite factorization.

Active-set constraint preconditioning

Suppose that the elements of D can be partitioned into two disjoint sets $\mathcal{B}$ and $\mathcal{S}$ that specify the "big" and "small" elements. Big elements tend to infinity. Small elements tend to zero.

We can define an *active-set preconditioner* as

$$(M + A_{\mathcal{S}}^T D_{\mathcal{S}}^{-1} A_{\mathcal{S}})$$

and

$$\begin{pmatrix} M + 2A_{\mathcal{S}}^T D_{\mathcal{S}}^{-1} A_{\mathcal{S}} & A_{\mathcal{S}}^T \\ A_{\mathcal{S}} & D_{\mathcal{S}} \end{pmatrix}$$

respectively.

Smaller preconditioning matrix with the same asymptotic properties as the constraint preconditioner.

Future research: Quasi-Newton preconditioners

The constraint preconditioners do not specify M very precisely.

Idea: Use limited memory quasi-Newton approximation of H to create a "good" M.

Understanding properties of the conjugate-gradient method important.

Research in progress.

Reference on interior methods

Reference:

- A. Forsgren, P. E. Gill and J. D. Griffin, *Iterative solution of augmented systems arising in interior methods*, Report TRITA-MAT-2005-OS3, Department of Mathematics, Royal Institute of Technology, 2005.

(Available from http://www.math.kth.se/~andersf)

Summary

The behavior of the conjugate-gradient method on ill-conditioned systems of linear equations have been discussed.

The behavior is beneficial in some application areas, e.g., radiation therapy optimization.

In other areas, suitable preconditioning is important for the efficiency of the method, e.g., interior methods for optimization.

Much interesting research ahead.

Thank you for your attention!

Stockholm, June 8 1990

The PhD student to be examined.

Stockholm, June 8 1990

The external examiner.

Stockholm, June 8 1990

The evaluation committee.