Numerical Optimization

Lecture Notes #10
Conjugate Gradient Methods — Linear CG, Part #1

Peter Blomgren,
⟨blomgren.peter@gmail.com⟩

Department of Mathematics and Statistics
Dynamical Systems Group
Computational Sciences Research Center
San Diego State University
San Diego, CA 92182-7720

http://terminus.sdsu.edu/

Fall 2018

Outline

Quick Recap: — Global Convergence and Enhancements

We looked at some theorems describing the convergence of our algorithms. We noted that there was a bit of a gap between what is generally true/practical, and what can be proved. (Theoretical limit points vs. numerical stopping criteria.)

Further, we looked at some enhancements including **scaling**

$$D = \text{diag}(d_1, d_2, \ldots, d_n), \quad d_i > 0, \quad T(\Delta) = \{\bar{\mathbf{p}} \in \mathbb{R}^n : \|D\bar{\mathbf{p}}\| \leq \Delta\},$$

and the use of **non-Euclidean norms** — the latter primarily come in handy in the context of constrained optimization.

We now explore an important computational tool, which will help us solve problems of realistic size. — **Conjugate Gradient Methods.**

Conjugate Gradient Methods: Introduction

For short: **"CG" Methods.**

- One of the most useful techniques for solving large linear systems of equations $A\bar{\mathbf{x}} = \bar{\mathbf{b}}$. **"Linear CG"**
- Can be adopted to solve nonlinear optimization problems. **"Nonlinear CG"** (Our type of problems!)
- Linear CG is an **alternative** to Gaussian elimination (well suited for large problems).
- Performance of linear CG is strongly tied to the distribution of the eigenvalues of A.

First, we explore the Linear CG method...

The Linear CG Method — Language and Notation

The **linear** CG method is an **iterative method** for solving linear systems of equations:

$$A\bar{\mathbf{x}} = \bar{\mathbf{b}}, \qquad A \in \mathbb{R}^{n\times n}, \quad \bar{\mathbf{x}} \in \mathbb{R}^n, \quad \bar{\mathbf{b}} \in \mathbb{R}^n,$$

where the matrix A is **symmetric positive definite**$^{\exists \text{ extensions}}$.

Notice/Recall: This problem is **equivalent to minimizing** $\Phi(\bar{\mathbf{x}})$ where

$$\Phi(\bar{\mathbf{x}}) = \frac{1}{2}\bar{\mathbf{x}}^T A\bar{\mathbf{x}} - \bar{\mathbf{b}}^T\bar{\mathbf{x}} + c,$$

since

$$\nabla\Phi(\bar{\mathbf{x}}) = A\bar{\mathbf{x}} - \bar{\mathbf{b}} \quad \stackrel{\text{def}}{=} \quad \bar{\mathbf{r}}(\bar{\mathbf{x}}).$$

We refer to $\bar{\mathbf{r}}(\bar{\mathbf{x}})$ as the **residual** of the linear system. Note that if $\bar{\mathbf{x}}^* = A^{-1}\bar{\mathbf{b}}$, then $\bar{\mathbf{r}}(\bar{\mathbf{x}}^*) = 0$, *i.e.* the residual is a measure of how close (or far) we are from solving the linear system.

Conjugate Directions

Definition (Conjugate Vector)

A set of nonzero vectors $\{\bar{\mathbf{p}}_0, \bar{\mathbf{p}}_1, \ldots, \bar{\mathbf{p}}_{n-1}\}$ is said to be **conjugate** with respect to the symmetric positive definite matrix A if

$$\bar{\mathbf{p}}_i^T A\bar{\mathbf{p}}_j = 0, \quad \forall i \neq j.$$

Property: Linear Independence of Conjugate Vectors

A set of conjugate vectors $\{\bar{\mathbf{p}}_0, \bar{\mathbf{p}}_1, \ldots, \bar{\mathbf{p}}_{n-1}\}$ is **linearly independent.**

Why should we care? — We can minimize $\Phi(\bar{\mathbf{x}})$ in n steps by successively minimizing along the directions in a conjugate set...

Conjugate Direction Method (!= CG Method) 1 of 4

Given a starting point $\bar{\mathbf{x}}_0 \in \mathbb{R}^n$, and a set of conjugate directions $\{\bar{\mathbf{p}}_0, \bar{\mathbf{p}}_1, \ldots, \bar{\mathbf{p}}_{n-1}\}$ we generate a sequence of points $\bar{\mathbf{x}}_k \in \mathbb{R}^n$ by setting

$$\bar{\mathbf{x}}_{k+1} = \bar{\mathbf{x}}_k + \alpha_k\bar{\mathbf{p}}_k,$$

where α_k is the minimizer of the quadratic function $\varphi(\alpha) = \Phi(\bar{\mathbf{x}}_k + \alpha\bar{\mathbf{p}}_k)$, *i.e.* the minimizer of $\Phi(\cdot)$ along the line $\bar{\ell}(\alpha) = \bar{\mathbf{x}}_k + \alpha\bar{\mathbf{p}}_k$.

We have already solved this problem — in the context of step-length selection for line search methods, **see** lecture #6 — so we "know" that the optimizer is given by

$$\alpha_k = -\frac{\bar{\mathbf{r}}_k^T\bar{\mathbf{p}}_k}{\bar{\mathbf{p}}_k^T A\bar{\mathbf{p}}_k}, \quad \text{where } \bar{\mathbf{r}}_k = \bar{\mathbf{r}}(\bar{\mathbf{x}}_k).$$

Conjugate Direction Method (!= CG Method) 2 of 4

Theorem (n-step convergence)

For any $\bar{\mathbf{x}}_0 \in \mathbb{R}^n$ the sequence $\{\bar{\mathbf{x}}_k\}$ generated by the conjugate direction algorithm converges to the solution $\bar{\mathbf{x}}^*$ of the linear system in at most n steps.

The proof indicates how properties of CG are found...

Proof: Part 1 (Fundmental Building Block).

Since the directions $\{\bar{\mathbf{p}}_i\}$ are linearly independent, they must **span** the whole space $\mathbb{R}^n$. Hence, we can write

$$\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_0 \quad = \quad \sum_{k=0}^{n-1}\sigma_k\bar{\mathbf{p}}_k$$

for some choice of scalars σ_k. We need to establish that $\sigma_k = \alpha_k$. □

Conjugate Direction Method (!= CG Method) 3 of 4

Proof: Part 2.

If we are generating $\bar{\mathbf{x}}_k$ by the conjugate direction method, then we have

$$\bar{\mathbf{x}}_k = \bar{\mathbf{x}}_0 + \alpha_0\bar{\mathbf{p}}_0 + \alpha_1\bar{\mathbf{p}}_1 + \cdots + \alpha_{k-1}\bar{\mathbf{p}}_{k-1},$$

we multiply this by $\bar{\mathbf{p}}_k^T A$

$$\bar{\mathbf{p}}_k^T A\bar{\mathbf{x}}_k = \bar{\mathbf{p}}_k^T A\left[\bar{\mathbf{x}}_0 + \alpha_0\bar{\mathbf{p}}_0 + \alpha_1\bar{\mathbf{p}}_1 + \cdots + \alpha_{k-1}\bar{\mathbf{p}}_{k-1}\right],$$

using the conjugacy property, we see that all but the first term on the right-hand-side are zero:

$$\bar{\mathbf{p}}_k^T A\bar{\mathbf{x}}_k = \bar{\mathbf{p}}_k^T A\bar{\mathbf{x}}_0 \quad \Leftrightarrow \quad \bar{\mathbf{p}}_k^T A(\bar{\mathbf{x}}_k - \bar{\mathbf{x}}_0) = 0.$$

Now we have

$$\bar{\mathbf{p}}_k^T A(\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_0) = \bar{\mathbf{p}}_k^T A(\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_0 - \underbrace{(\bar{\mathbf{x}}_k - \bar{\mathbf{x}}_0)}_{\text{adds 0}}) = \bar{\mathbf{p}}_k^T A(\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_k) = \bar{\mathbf{p}}_k^T(\bar{\mathbf{b}} - A\bar{\mathbf{x}}_k) = -\bar{\mathbf{p}}_k^T\bar{\mathbf{r}}_k.$$

□

Conjugate Direction Method (!= CG Method) 4 of 4

Proof: Part 3.

We have shown

$$\bar{\mathbf{p}}_k^T A(\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_0) = -\bar{\mathbf{p}}_k^T\bar{\mathbf{r}}_k.$$

Now, we notice that the right-hand-side is the numerator in α_k:

$$\alpha_k = \frac{-\bar{\mathbf{p}}_k^T\bar{\mathbf{r}}_k}{\bar{\mathbf{p}}_k^T A\bar{\mathbf{p}}_k} \quad \Rightarrow \quad \alpha_k = \frac{\bar{\mathbf{p}}_k^T A(\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_0)}{\bar{\mathbf{p}}_k^T A\bar{\mathbf{p}}_k}.$$

We conclude the proof by showing that σ_k can be expressed in the same manner; we premultiply the expression for $(\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_0)$ by $\bar{\mathbf{p}}_k^T A$ and obtain

$$\bar{\mathbf{p}}_k^T A(\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_0) = \bar{\mathbf{p}}_k^T A\sum_{i=0}^{n-1}\sigma_i\bar{\mathbf{p}}_i = \sum_{i=0}^{n-1}\sigma_i\bar{\mathbf{p}}_k^T A\bar{\mathbf{p}}_i = \sigma_k\bar{\mathbf{p}}_k^T A\bar{\mathbf{p}}_k.$$

Hence,

$$\sigma_k = \frac{\bar{\mathbf{p}}_k^T A(\bar{\mathbf{x}}^* - \bar{\mathbf{x}}_0)}{\bar{\mathbf{p}}_k^T A\bar{\mathbf{p}}_k} \quad \equiv \quad \alpha_k.$$

□

Conjugate Direction Method: Comments and Interpretation 1 of 2

Most of the proofs regarding CD and CG methods are argued in a similar way — by looking at optimizers and residuals over sub-spaces of $\mathbb{R}^n$ spanned by some subset of a set of conjugate vectors.

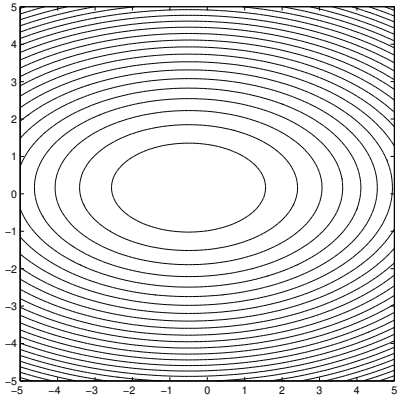

Interpretation: If the matrix A is diagonal, then the contours of $\Phi(\bar{\mathbf{x}})$ are ellipses whose axes are aligned with the coordinate directions. In this case, we can find the minimizer by performing 1D-minimizations along the coordinate directions $\bar{\mathbf{e}}_1, \bar{\mathbf{e}}_2, \ldots, \bar{\mathbf{e}}_n$ in turn.

Conjugate Direction Method: Comments and Interpretation 2 of 2

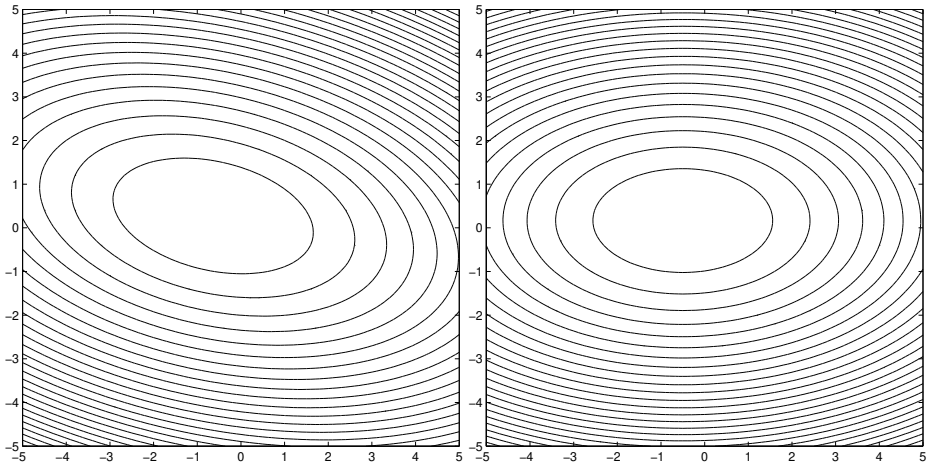

Interpretation (ctd.): When A is not diagonal, the contours are still elliptical, but are no longer aligned with the coordinate axes. Successive minimization along the coordinate directions $\bar{\mathbf{e}}_1, \bar{\mathbf{e}}_2, \ldots, \bar{\mathbf{e}}_n$ can **not** guarantee convergence in n (or even a (fixed) finite number of) iterations.

Recovering n-step Convergence for Non-Diagonal A 1 of 2

For non-diagonal matrices A, the n-step convergence can be recovered by transforming the problem.

Let $S \in \mathbb{R}^{n \times n}$ be a matrix with conjugate columns, *i.e.* if $\{\bar{\mathbf{p}}_0, \bar{\mathbf{p}}_1, \ldots, \bar{\mathbf{p}}_{n-1}\}$ is a set of conjugate directions (with respect to A), then

$$S = \begin{bmatrix} | & | & & | \\ \bar{\mathbf{p}}_0 & \bar{\mathbf{p}}_1 & \cdots & \bar{\mathbf{p}}_{n-1} \\ | & | & & | \end{bmatrix}.$$

We introduce a new variable $\widehat{\mathbf{x}} = S^{-1}\bar{\mathbf{x}}$, and thus get the new quadratic objective which can be minimized in n steps

$$\widehat{\Phi}(\widehat{\mathbf{x}}) = \Phi(S\widehat{\mathbf{x}}) = \frac{1}{2}\widehat{\mathbf{x}}^T \underbrace{(S^T A S)}_{\text{Diagonal}} \widehat{\mathbf{x}} - (S^T \bar{\mathbf{b}})^T \widehat{\mathbf{x}}.$$

Recovering n-step Convergence for Non-Diagonal A 2 of 2

We note that the matrix $(S^T A S)$ is diagonal by the conjugacy property, and that each coordinate direction $\widehat{\mathbf{e}}_i$ in $\widehat{\mathbf{x}}$-space corresponds to the direction $\bar{\mathbf{p}}_{i-1}$ in $\bar{\mathbf{x}}$-space.

When the matrix is diagonal, each coordinate minimization determines one of the components of the solution $\bar{\mathbf{x}}^*$. Hence, after k iterations, the quadratic has been minimized on the subspace spanned by $\widehat{\mathbf{e}}_1, \widehat{\mathbf{e}}_2, \ldots, \widehat{\mathbf{e}}_k$.

If we instead minimize along the conjugate directions, then after k iterations, the quadratic has been minimized on the subspace spanned by $\bar{\mathbf{p}}_0, \bar{\mathbf{p}}_1, \ldots, \bar{\mathbf{p}}_{k-1}$.

Updating the Residual

Before we state a fundamental theorem regarding the conjugate direction method, we show the following lemma:

Lemma

Given a starting point $\bar{\mathbf{x}}_0 \in \mathbb{R}^n$ and a set of conjugate directions $\{\bar{\mathbf{p}}_0, \bar{\mathbf{p}}_1, \ldots, \bar{\mathbf{p}}_{n-1}\}$ if we generate the sequence $\bar{\mathbf{x}}_k \in \mathbb{R}^n$ by setting

$$\bar{\mathbf{x}}_{k+1} = \bar{\mathbf{x}}_k + \alpha_k \bar{\mathbf{p}}_k, \quad \text{where } \alpha_k = -\frac{\bar{\mathbf{r}}_k^T \bar{\mathbf{p}}_k}{\bar{\mathbf{p}}_k^T A \bar{\mathbf{p}}_k},$$

with $\bar{\mathbf{r}}_k = A\bar{\mathbf{x}}_k - b$. Then the $(k+1)$st residual is given by the following expression

$$\bar{\mathbf{r}}_{k+1} = \bar{\mathbf{r}}_k + \alpha_k A \bar{\mathbf{p}}_k.$$

Proof: (Quick One-Liner).

$$\bar{\mathbf{r}}_{k+1} = A\bar{\mathbf{x}}_{k+1} - \bar{\mathbf{b}} = A(\bar{\mathbf{x}}_k + \alpha_k \bar{\mathbf{p}}_k) - \bar{\mathbf{b}} = \alpha_k A \bar{\mathbf{p}}_k + (A\bar{\mathbf{x}}_k - \bar{\mathbf{b}}) = \alpha_k A \bar{\mathbf{p}}_k + \bar{\mathbf{r}}_k.$$

□

Expanding Subspace Minimization

Theorem (Expanding Subspace Minimization)

Let $\bar{\mathbf{x}}_0 \in \mathbb{R}^n$ be any starting point and suppose that the sequence $\{\bar{\mathbf{x}}_k\}$ is generated by

$$\bar{\mathbf{x}}_{k+1} = \bar{\mathbf{x}}_k + \alpha_k \bar{\mathbf{p}}_k, \quad \text{where } \alpha_k = -\frac{\bar{\mathbf{r}}_k^T \bar{\mathbf{p}}_k}{\bar{\mathbf{p}}_k^T A \bar{\mathbf{p}}_k}.$$

Then

$$\bar{\mathbf{r}}_k^T \bar{\mathbf{p}}_i = 0, \quad \text{for } i = 0, 1, \ldots, k-1,$$

and $\bar{\mathbf{x}}_k$ is the minimizer of $\Phi(\bar{\mathbf{x}}) = \frac{1}{2}\bar{\mathbf{x}}^T A \bar{\mathbf{x}} - \bar{\mathbf{b}}^T \bar{\mathbf{x}}$ over the set

$$S(\bar{\mathbf{x}}_0, k) = \left\{ \bar{\mathbf{x}} \,:\, \bar{\mathbf{x}} = \bar{\mathbf{x}}_0 + \operatorname{span}\{\bar{\mathbf{p}}_0, \bar{\mathbf{p}}_1, \ldots, \bar{\mathbf{p}}_{k-1}\} \right\}.$$

Expanding Subspace Minimization: Proof 1 of 3

Proof: Part 1 (Fundmental Building Block).

First, we show that a point $\tilde{\mathbf{x}}$ minimizes Φ over the set $S(\bar{\mathbf{x}}_0, k)$ if and only if $r(\tilde{\mathbf{x}})^T \bar{\mathbf{p}}_i = 0,\ i = 0, 1, \ldots, k-1$.

Let $h(\bar{\sigma}) = \Phi(\bar{\mathbf{x}}_0 + \sigma_0 \bar{\mathbf{p}}_0 + \sigma_1 \bar{\mathbf{p}}_1 + \cdots + \sigma_{k-1} \bar{\mathbf{p}}_{k-1})$. Since $h(\bar{\sigma})$ is a strictly convex quadratic it has a unique minimizer $\bar{\sigma}^*$ that satisfies

$$\frac{\partial h(\bar{\sigma}^*)}{\partial \sigma_i} = 0, \quad i = 0, 1, \ldots, k-1$$

By the chain rule, this is equivalent to

$$\nabla \Phi(\underbrace{\bar{\mathbf{x}}_0 + \sigma_0^* \bar{\mathbf{p}}_0 + \sigma_1^* \bar{\mathbf{p}}_1 + \cdots + \sigma_{k-1}^* \bar{\mathbf{p}}_{k-1}}_{\tilde{\mathbf{x}}})^T \bar{\mathbf{p}}_i = 0, \quad i = 0, 1, \ldots, k-1$$

We recall that $\nabla \Phi(\tilde{\mathbf{x}}) = A\tilde{\mathbf{x}} - \bar{\mathbf{b}} = \bar{\mathbf{r}}(\tilde{\mathbf{x}})$, thus we have established $\bar{\mathbf{r}}(\tilde{\mathbf{x}})^T \bar{\mathbf{p}}_i = 0 \Leftrightarrow \tilde{\mathbf{x}}$ minimizes Φ over the set $S(\bar{\mathbf{x}}_0, k)$. □

Expanding Subspace Minimization: Proof 2 of 3

Proof: Part 2.

We now show that the residuals $\bar{\mathbf{r}}_k$ satisfy $\bar{\mathbf{r}}_k^T \bar{\mathbf{p}}_i = 0,\ i = 0, 1, \ldots, k-1$.

We use mathematical induction. Since α_0 is always the 1D-minimizer, we have $\bar{\mathbf{r}}_1^T \bar{\mathbf{p}}_0 = 0$, establishing the **base case**.

From the **inductive hypothesis**, that $\bar{\mathbf{r}}_{k-1}^T \bar{\mathbf{p}}_i = 0,\ i = 0, 1, \ldots, k-2$, we must show that $\bar{\mathbf{r}}_k^T \bar{\mathbf{p}}_i = 0,\ i = 0, 1, \ldots, k-1$ in order to complete the proof.

From the lemma we have an expression for $\bar{\mathbf{r}}_k = \bar{\mathbf{r}}_{k-1} + \alpha_{k-1} A \bar{\mathbf{p}}_{k-1}$.

First off we have: $\bar{\mathbf{p}}_{k-1}^T \bar{\mathbf{r}}_k = \bar{\mathbf{p}}_{k-1}^T \bar{\mathbf{r}}_{k-1} + \alpha_{k-1} \bar{\mathbf{p}}_{k-1}^T A \bar{\mathbf{p}}_{k-1} = 0$, since, by construction (optimality)

$$\alpha_{k-1} = \frac{-\bar{\mathbf{p}}_{k-1}^T \bar{\mathbf{r}}_{k-1}}{\bar{\mathbf{p}}_{k-1}^T A \bar{\mathbf{p}}_{k-1}}$$

□

Expanding Subspace Minimization: Proof 3 of 3

Proof: Part 3.

Finally,

$$\bar{\mathbf{p}}_i^T \bar{\mathbf{r}}_k = \bar{\mathbf{p}}_i^T \bar{\mathbf{r}}_{k-1} + \alpha_{k-1} \bar{\mathbf{p}}_i^T A \bar{\mathbf{p}}_{k-1} = 0, \quad i = 0, 1, \ldots, k-2$$

since

$$\bar{\mathbf{p}}_i^T \bar{\mathbf{r}}_{k-1} = 0, \quad i = 0, 1, \ldots, k-2$$

by the induction hypothesis, and

$$\bar{\mathbf{p}}_i^T A \bar{\mathbf{p}}_{k-1} = 0, \quad i = 0, 1, \ldots, k-2$$

by conjugacy. This establishes $\bar{\mathbf{p}}_i^T \bar{\mathbf{r}}_k = 0,\ i = 0, 1, \ldots, k-1$, which completes the proof. □

Cliff-Hangers...

Cliff-Hanger Questions:

- How can we make this useful?
- Given A, how do we get a set of conjugate vectors? (They are not for sale at Costco!)
- Even if we have them, why is this scheme any better than Gaussian elimination?
- Where is the gradient?

Index