Lecture 10: CS395T Numerical Optimization for Graphics and AI — Conjugate Gradient Methods (Linear)

Qixing Huang
The University of Texas at Austin
huangqx@cs.utexas.edu

1 Disclaimer

This note is adapted from

- Section 5 of *Numerical Optimization* by Jorge Nocedal and Stephen J. Wright. Springer series in operations research and financial engineering. Springer, New York, NY, 2. ed. edition, (2006)

2 Basic Setting

In this section, we study how to solve the following optimization problem using conjugate gradient method

$$\underset{\boldsymbol{x}}{\text{minimize}} \quad -\boldsymbol{b}^T\boldsymbol{x} + \frac{1}{2}\boldsymbol{x}^T A\boldsymbol{x},$$

where A is assumed to be positive definite. Later we will replace this constraint and also consider non-linear objective functions.

Conjugate directions. We say a few directions $\boldsymbol{p}_i, 1 \leq i \leq n$ are conjugate with respect to A if

$$\boldsymbol{p}_i^T A\boldsymbol{p}_j = 0, \qquad 1 \leq i \neq j \leq n.$$

It is easy to see that when $A = I_n$, then conjugate directions become orthogonal directions. This means conjugate directions are essentially generalizations of orthogonal directions.

Conjugate directions also exist, for example, let the spectral decomposition of A be

$$A = U\Sigma U^T.$$

Then we can let $\boldsymbol{p}_1, \cdots, \boldsymbol{p}_n$ be the columns of U, i.e., $U = (\boldsymbol{p}_1, \cdots, \boldsymbol{p}_n)$, and it is easy to check that they are indeed conjugate directions. On the other hand, the set of conjugate directions are not unique.

Fact 2.1. *Conjugate directions are linearly independent.*

Proof. Suppose they are linearly dependent. This means there exist non-zero (meaning not all of them are zero) coefficients $c_i, 1 \leq i \leq n$ such that

$$\sum_{i=1}^{n} c_i\boldsymbol{p}_i = 0.$$

It follows that

$$
\begin{aligned}
0 &= (\sum_{i=1}^{n} c_i \boldsymbol{p}_i)^T A (\sum_{i=1}^{n} c_i \boldsymbol{p}_i) \\
&= \sum_{i=1}^{n} c_i^2 (\boldsymbol{p}_i^T A \boldsymbol{p}_i) > 0,
\end{aligned} \tag{1}
$$

which leads to a contradiction. □

The conjugate directions are useful for optimization. Intuitively, in the case $A = I_n$, we can essentially optimize along each coordinate independently to obtain the optimal solution. This is also true in the general case. In fact, starting from an initial solution $\boldsymbol{x}_0$, we can gradually improve the solution by searching along $\boldsymbol{p}_k$ at each iteration:

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

where

$$
\begin{aligned}
\alpha_k &= \underset{\alpha}{\text{minimize}}\ \boldsymbol{b}^T (\boldsymbol{x}_{k-1} + \alpha \boldsymbol{p}_k) - \frac{1}{2} (\boldsymbol{x}_{k-1} + \alpha \boldsymbol{p}_k)^T A (\boldsymbol{x}_{k-1} + \alpha \boldsymbol{p}_k) \\
&= -\frac{1}{2} \alpha^2 (\boldsymbol{p}_k^T A \boldsymbol{p}_k) + \alpha (\boldsymbol{b} - A \boldsymbol{x}_k)^T \boldsymbol{p}_k \\
&= -\frac{\boldsymbol{r}_k^T \boldsymbol{p}_k}{\boldsymbol{p}_k^T A \boldsymbol{p}_k},
\end{aligned} \tag{2}
$$

where $\boldsymbol{r}_k := \boldsymbol{b} - A\boldsymbol{x}_k$.

Proposition 1. *The procedure described above converges to the optimal solution $\boldsymbol{x}^\star = -A^{-1}\boldsymbol{b}$ in at most n steps.*

Proof. Define

$$
\boldsymbol{r}_k := A\boldsymbol{p}_k - \boldsymbol{b}.
$$

It is equivalent to show that

$$
\boldsymbol{r}_k^T \boldsymbol{p}_i = 0, \quad 1 \le i < k \le n.
$$

We prove this by induction. When $k = 1$, we have

$$
\begin{aligned}
\boldsymbol{r}_1^T \boldsymbol{p}_0 &= -(\boldsymbol{b} - A\boldsymbol{x}_0 - \alpha_0 A \boldsymbol{p}_0)^T \boldsymbol{p}_1 \\
&= \boldsymbol{r}_0^T \boldsymbol{p}_0 - \alpha_0 \cdot (\boldsymbol{p}_0^T A \boldsymbol{p}_0) = 0.
\end{aligned}
$$

Suppose it is true for $1 \le i < k \le j$, now let us consider $k = j + 1$. First of all, we have

$$
\begin{aligned}
\boldsymbol{r}_{j+1}^T \boldsymbol{p}_j &= -(\boldsymbol{b} - A\boldsymbol{x}_j - \alpha_j A \boldsymbol{p}_j)^T \boldsymbol{p}_j \\
&= \boldsymbol{r}_j^T \boldsymbol{p}_j - \alpha_j (\boldsymbol{p}_j^T A \boldsymbol{p}_j) = 0.
\end{aligned}
$$

When $i < j + 1$, we have

$$
\begin{aligned}
\boldsymbol{r}_{j+1}^T \boldsymbol{p}_i &= -(\boldsymbol{b} - A\boldsymbol{x}_j - \alpha_j A \boldsymbol{p}_j)^T \boldsymbol{p}_i \\
&= -\boldsymbol{r}_j^T \boldsymbol{p}_i + \alpha_j \boldsymbol{p}_j^T A \boldsymbol{p}_i = 0,
\end{aligned}
$$

which ends the proof. □

The conjugate gradient method is based on a genius idea of choosing the conjugate gradient directions. Specifically, let

$$
\boldsymbol{p}_0 = -\boldsymbol{r}_0 = \boldsymbol{b} - A\boldsymbol{x}_0.
$$

The intermediate search directions are given by

$$
\boldsymbol{p}_k = -\boldsymbol{r}_k + \beta_k \boldsymbol{p}_{k-1},
$$

where β_k is chosen such that $\boldsymbol{p}_k^T A \boldsymbol{p}_{k-1} = 0$ or in other words

$$\beta_k = \frac{\boldsymbol{r}_k^T A \boldsymbol{p}_{k-1}}{\boldsymbol{p}_{k-1}^T A \boldsymbol{p}_{k-1}}.$$

Interesting, we can show that $\boldsymbol{p}_k$ is also conjugate to previous directions. Formally speaking, we summarize the properties in the following theorem:

Theorem 2.1. *Suppose that the kth iterate generated by the conjugate gradient method is not the solution point $\boldsymbol{x}^\star$. The following four properties hold:*

$$\boldsymbol{r}_k^T \boldsymbol{r}_i = 0, \qquad \text{for } i = 0, 1, \ldots, k-1, \tag{3}$$

$$\text{span}\{\boldsymbol{r}_0, \boldsymbol{r}_1, \cdots, \boldsymbol{r}_k\} = \text{span}\{\boldsymbol{r}_0, A\boldsymbol{r}_0, \ldots, A^k \boldsymbol{r}_0\}, \tag{4}$$

$$\text{span}\{\boldsymbol{p}_0, \boldsymbol{p}_1, \cdots, \boldsymbol{p}_k\} = \text{span}\{\boldsymbol{r}_0, A\boldsymbol{r}_0, \ldots, A^k \boldsymbol{r}_0\}, \tag{5}$$

$$\boldsymbol{p}_k^T A \boldsymbol{p}_i = 0, \qquad \text{for } i = 0, 1, \ldots, k-1. \tag{6}$$

Therefore, the sequence $\{\boldsymbol{x}_k\}$ converges to $\boldsymbol{x}^\star$ in at most n steps.

Proof. The proof is by induction. The expressions (4) and (5) holds trivially when $k = 0$. In addition, by construction $\boldsymbol{p}_1^T A \boldsymbol{p}_0 = 0$ so (6) also holds when $k = 1$. Now suppose (3)-(6) hold for all $j \leq k$. Now consider $k+1$. Since $\boldsymbol{p}_0, \cdots, \boldsymbol{p}_k$ are conjugate directions, we conclude that

$$\boldsymbol{r}_k^T \boldsymbol{p}_i = 0, \qquad 0 \leq i \leq k.$$

We first show that

$$\alpha_k \neq 0.$$

Suppose $\alpha_k = 0$, this means

$$\begin{aligned} 0 &= \boldsymbol{r}_k^T A \boldsymbol{p}_k \\ &= \boldsymbol{r}_k^T A(-\boldsymbol{r}_k + \beta_k \boldsymbol{p}_{k-1}) \\ &= -\boldsymbol{r}_k^T A \boldsymbol{r}_k + \frac{1}{\alpha_{k-1}} \boldsymbol{r}_k^T (\boldsymbol{r}_{k-1} - \boldsymbol{r}_{k-2}) \\ &= -\boldsymbol{r}_k^T A \boldsymbol{r}_k. \end{aligned}$$

Since A is positive semidefinite, it means $\boldsymbol{r}_k = 0$, which leads to a contradiction.

We then consider (4) and (5). Since $\boldsymbol{r}_{k+1} = \alpha_k A \boldsymbol{p}_k + \boldsymbol{r}_k$. As by induction $\boldsymbol{p}_k \in \text{span}\{\boldsymbol{r}_0, A\boldsymbol{r}_0, \ldots, A^k \boldsymbol{r}_0\}$. It follows that $\boldsymbol{r}_{k+1} \in \text{span}\{\boldsymbol{r}_0, A\boldsymbol{r}_0, \ldots, A^{k+1}\boldsymbol{r}_0\}$. This means

$$\text{span}\{\boldsymbol{r}_0, \boldsymbol{r}_1, \cdots, \boldsymbol{r}_{k+1}\} \subset \text{span}\{\boldsymbol{r}_0, A\boldsymbol{r}_0, \ldots, A^{k+1}\boldsymbol{r}_0\}.$$

In addition, since $\boldsymbol{p}_{k+1} = -\boldsymbol{r}_{k+1} + \beta_{k+1}\boldsymbol{p}_k$, it follows that

$$\text{span}\{\boldsymbol{p}_0, \boldsymbol{r}_1, \cdots, \boldsymbol{p}_{k+1}\} \subset \text{span}\{\boldsymbol{r}_0, A\boldsymbol{r}_0, \ldots, A^{k+1}\boldsymbol{r}_0\}.$$

To prove the inverse. Let

$$A^k \boldsymbol{r}_0 = c_0 \boldsymbol{r}_0 + \cdots + c_k \boldsymbol{r}_k.$$

It follows that

$$
\begin{aligned}
A^{k+1}\boldsymbol{r}_0 &\in \mathrm{span}\{A\boldsymbol{r}_0, \cdots, A\boldsymbol{r}_k\} \\
&\in \mathrm{span}\{A\boldsymbol{r}_0, A(\beta_1\boldsymbol{p}_0 - \boldsymbol{p}_1), \cdots, A(\beta_k\boldsymbol{p}_{k-1} - \boldsymbol{p}_k)\} \\
&\in \mathrm{span}\{A\boldsymbol{p}_0, A\boldsymbol{p}_1, \cdots, A\boldsymbol{p}_k\} \\
&\in \mathrm{span}\{\frac{(\boldsymbol{r}_1 - \boldsymbol{r}_0)}{\alpha_0}, \frac{(\boldsymbol{r}_2 - \boldsymbol{r}_1)}{\alpha_1}, \cdots, \frac{(\boldsymbol{r}_{k+1} - \boldsymbol{r}_k)}{\alpha_k}\} \\
&\in \mathrm{span}\{\boldsymbol{r}_0, \cdots, \boldsymbol{r}_{k+1}\} \\
&\in \mathrm{span}\{\boldsymbol{p}_0, \beta_1\boldsymbol{p}_0 - \boldsymbol{p}_1 \cdots, \beta_{k+1}\boldsymbol{p}_k - \boldsymbol{p}_{k+1}\} \\
&\in \mathrm{span}\{\boldsymbol{p}_0, \boldsymbol{p}_1 \cdots, \boldsymbol{p}_{k+1}\}
\end{aligned}
$$

Now we show that

$$\boldsymbol{r}_{k+1}^T\boldsymbol{r}_i = 0, \qquad 0 \le i \le k.$$

In fact,

$$\boldsymbol{r}_{k+1}^T\boldsymbol{r}_i = \boldsymbol{r}_{k+1}^T(\beta_i\boldsymbol{p}_{i-1} - \boldsymbol{p}_i) = \beta_i(\boldsymbol{r}_{k+1}^T\boldsymbol{p}_{i-1}) - \boldsymbol{r}_{k+1}^T\boldsymbol{p}_i = 0.$$

Finally, for $0 \le i \le k$,

$$\boldsymbol{p}_{k+1}^T A\boldsymbol{p}_i = (-\boldsymbol{r}_{k+1} + \beta_{k+1}\boldsymbol{p}_k)^T A\boldsymbol{p}_i = -\boldsymbol{r}_{k+1}^T A\boldsymbol{p}_i = -\boldsymbol{r}_{k+1}^T(\frac{\boldsymbol{r}_i - \boldsymbol{r}_{i-1}}{\alpha_i}) = 0.$$

□

Conjugate gradient methods require computing two constants

$$\beta_k = \frac{\boldsymbol{r}_{k+1}^T Ap_k}{\boldsymbol{p}_k^T A\boldsymbol{p}_k}, \quad \alpha_k = -\frac{\boldsymbol{r}_k^T\boldsymbol{p}_k}{\boldsymbol{p}_k^T A\boldsymbol{p}_k}.$$

Note that

$$\boldsymbol{r}_k^T\boldsymbol{p}_k = -\boldsymbol{r}_k^T(-\boldsymbol{r}_k + \beta_k\boldsymbol{p}_{k-1}) = \boldsymbol{r}_k^T\boldsymbol{r}_k.$$

Moreover,

$$
\begin{aligned}
\beta_k &= \frac{\boldsymbol{r}_{k+1}^T A\boldsymbol{p}_k}{\boldsymbol{p}_k^T A\boldsymbol{p}_k} = \frac{\boldsymbol{r}_{k+1}^T A\boldsymbol{p}_k}{\boldsymbol{p}_k^T A(-\boldsymbol{r}_k + \beta_k\boldsymbol{p}_{k-1})} = -\frac{\boldsymbol{r}_{k+1}^T A\boldsymbol{p}_k}{\boldsymbol{r}_k^T A\boldsymbol{p}_k} = -\frac{\boldsymbol{r}_{k+1}^T(\alpha_k A\boldsymbol{p}_k)}{\boldsymbol{r}_k^T(\alpha_k A\boldsymbol{p}_k)} \\
&= -\frac{\boldsymbol{r}_{k+1}^T(\boldsymbol{r}_{k+1} - \boldsymbol{r}_k)}{\boldsymbol{r}_k^T(\boldsymbol{r}_{k+1} - \boldsymbol{r}_k)} = \frac{\boldsymbol{r}_{k+1}^T\boldsymbol{r}_{k+1}}{\boldsymbol{r}_k^T\boldsymbol{r}_k}.
\end{aligned}
$$

Finally, we arrive at the standard form of conjugate gradient descent method:

- Given $\boldsymbol{x}_0$;
- Set $\boldsymbol{r}_0 = A\boldsymbol{x}_0 - \boldsymbol{b}, \boldsymbol{p}_0 \leftarrow -\boldsymbol{r}_0, k \leftarrow 0$;
- while $r_k \neq 0$
- $\alpha_k \leftarrow \frac{\boldsymbol{r}_k^T\boldsymbol{r}_k}{\boldsymbol{p}_k^T A\boldsymbol{p}_k}$;
- $\boldsymbol{x}_{k+1} \leftarrow \boldsymbol{x}_k + \alpha_k\boldsymbol{p}_k$;
- $\boldsymbol{r}_{k+1} \leftarrow \boldsymbol{r}_k + \alpha_k A\boldsymbol{p}_k$;
- $\beta_{k+1} \leftarrow \frac{\boldsymbol{r}_{k+1}^T\boldsymbol{r}_{k+1}}{\boldsymbol{r}_k^T\boldsymbol{r}_k}$;
- $\boldsymbol{p}_{k+1} = -\boldsymbol{r}_{k+1}^T + \beta_{k+1}\boldsymbol{p}_k$;
- $k \leftarrow k + 1$;
- end (while)

2.1 Convergence Rate of CG

One thing is analyze the convergence of CG is to utilize the fact that $\boldsymbol{x}_{k+1}$ is the approximation of $\boldsymbol{x}^\star$ in $\boldsymbol{x}_0 + \mathrm{span}(\boldsymbol{r}_0, A\boldsymbol{r}_0, \cdots, A^k\boldsymbol{r}_0)$. Define

$$P_k(A) = \gamma_0 I + \gamma_1 A + \cdots + \gamma_k A^k,$$

where $\gamma_0, \gamma_1, \cdots, \gamma_k$ are coefficients. It turns out $\boldsymbol{x}_{k+1}$ is given by the best degree k polynomial that solves the following optimization problem:

$$\min_{P_k} \|\boldsymbol{x}_0 + P_k(A)\boldsymbol{r}_0 - \boldsymbol{x}^\star\|_A^2.$$

Let $P_k^\star(A)$ be the best polynomial, we have that

$$\boldsymbol{x}_{k+1} - \boldsymbol{x}^\star = \boldsymbol{x}_0 + P_k^\star(A)\boldsymbol{r}_0 - \boldsymbol{x}^\star = (I + P_k^\star(A)A)(\boldsymbol{x}_0 - \boldsymbol{x}^\star).$$

Let $\boldsymbol{v}_i, 1 \le i \le n$ be the eigen-vectors of A. Write

$$\boldsymbol{x}_0 - \boldsymbol{x}^\star = \sum_{i=1}^n \psi_i \boldsymbol{v}_i.$$

It follows that

$$\|\boldsymbol{x}_{k+1} - \boldsymbol{x}^\star\|_A^2 = \sum_{i=1}^n \lambda_i (1 + \lambda_i P_k^\star(\lambda_i))\psi_i^2.$$

So we have

$$\|\boldsymbol{x}_{k+1} - \boldsymbol{x}^\star\|_A^2 = \min_{P_k} \sum_{i=1}^n \lambda_i (1 + \lambda_i P_k(\lambda_i))\psi_i^2.$$

It follows that

$$\|\boldsymbol{x}_{k+1} - \boldsymbol{x}^\star\|_A^2 \le \min_{P_k} \max_{1 \le i \le n} (1 + \lambda_i P_k(\lambda_i))^2 \|\boldsymbol{x}_0 - \boldsymbol{x}^\star\|_A^2.$$

By choosing different polynomials, we can have different types convergence bounds. In the literature, people have obtained the following bounds:

$$\|\boldsymbol{x}_{k+1} - \boldsymbol{x}^\star\|_A^2 \le \left(\frac{\lambda_{k+1} - \lambda_n}{\lambda_{k+1} + \lambda_n}\right)^2 \|\boldsymbol{x}_0 - \boldsymbol{x}^\star\|_A^2 \tag{7}$$

and

$$\|\boldsymbol{x}_{k+1} - \boldsymbol{x}^\star\|_A^2 \le \left(\frac{\sqrt{\kappa(A)} - 1}{\sqrt{\kappa(A)} + 1}\right)^{k+1} \|\boldsymbol{x}_0 - \boldsymbol{x}^\star\|_A^2. \tag{8}$$