Lecture 16: CS395T Numerical Optimization for Graphics and AI — Linear Programming (Simplex Method II and Interior Point Method I

Qixing Huang

The University of Texas at Austin

`huangqx@cs.utexas.edu`

Disclaimer

This note is adapted from

- Section 13 and Section 14 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)

1 Simplex Methods

In class we will discuss the following items regarding the simplex method:

Dual Program and Optimality Conditions. The primal problem of LP

$$\min_{\boldsymbol{x} \in \mathbb{R}^n} \boldsymbol{c}^T \boldsymbol{x}$$
$$\text{subject to } A\boldsymbol{x} = \boldsymbol{b}$$
$$\boldsymbol{x} \geq 0. \tag{1}$$

The dual problem of LP

$$\max_{\lambda, \boldsymbol{s}} \boldsymbol{b}^T \lambda$$
$$\text{subject to } A^T \lambda + \boldsymbol{s} = \boldsymbol{c}, \ \boldsymbol{s} \geq 0$$
$$\boldsymbol{x} \geq 0. \tag{2}$$

An important property of LP is that the KKT conditions are sufficient for optimality:

$$A^T \lambda + \boldsymbol{s} = \boldsymbol{c}, \tag{3}$$
$$A\boldsymbol{x} = \boldsymbol{b}, \tag{4}$$
$$\boldsymbol{x} \geq 0, \tag{5}$$
$$\boldsymbol{s} \geq 0, \tag{6}$$
$$x_i s_i = 0, \ i = 1, 2, \cdots, n. \tag{7}$$

Geometry of Feasible Set.

Definition 1.1. *A vector $\boldsymbol{x}$ is a basic feasible point if it is feasible and if there exists a subset $\mathcal{B}$ of the index set $\{1, \cdots, n\}$ such that*

- $\mathcal{B}$ contains exactly m indices;

- $i \notin \mathcal{B} \rightarrow x_i = 0$ (that is, the bound $x_i \geq 0$ can be inactive only if $i \in \mathcal{B}$);

- The $m \times m$ matrix B defined by

$$B = [A_i]_{i \in \mathcal{B}}$$

  is non-singular, where A_i is the i-th column of A.

A set $\mathcal{B}$ satisfying these properties is called a basis for the problem (1). The corresponding matrix B is called the basis matrix.

Theorem 1.1.
- If (1) has a nonempty feasible region, then there is at least one basic feasible point;

- If (1) has solutions, then at least one such solution is a basic optimal point.

- If (1) is feasible and bounded, then it has an optimal solution.

Theorem 1.2. All basic feasible points for (1) are vertices of the feasible polytope $\{\boldsymbol{x}|A\boldsymbol{x} = \boldsymbol{b}, \boldsymbol{x} \geq 0\}$, and vice versa.

Definition 1.2. A basis $\mathcal{B}$ is said to be degenerate if $x_i = 0$ for some $i \in \mathcal{B}$, where $\boldsymbol{x}$ is the basic feasible solution corresponding to $\mathcal{B}$. A linear program is said to be degenerate if it has at least one degenerate basis.

Two-Phase Procedure for the Simplex Method. The first phase solves the following linear program to obtain an initial solution:

$$\min \boldsymbol{e}^T \boldsymbol{z} \quad \text{subject to} \quad A\boldsymbol{x} + E\boldsymbol{z} = \boldsymbol{b}, \ (\boldsymbol{x}, \boldsymbol{z}) \geq 0, \tag{8}$$

where $\boldsymbol{z} \in \mathbb{R}^m$, $\boldsymbol{e} = (1, \cdots, 1)^T$, and E is a diagonal matrix whose diagonal elements are

$$E_{jj} = 1, \text{ if } b_j \geq 0, \quad E_{jj} = -1, \text{ if } b_j = 0.$$

The nice thing about this formulation is that there exists a very simple basic feasible solution for (8):

$$\boldsymbol{x} = 0, \quad z_j = |b_j|, j = 1, 2, \cdots, m.$$

The second phase solves the following linear program:

$$\min \boldsymbol{c}^T \boldsymbol{x} \quad \text{subject to} \quad A\boldsymbol{x} + \boldsymbol{z} = \boldsymbol{b}, \ \boldsymbol{x} \geq 0, \quad 0 \geq \boldsymbol{z} \quad 0. \tag{9}$$

It is easy to modify the simplex method for solving (9).

2 Interior Point Method (Primal-Dual Methods)

2.1 Outline

Primal-dual methods find solutions $(\boldsymbol{x}^\star, \lambda^\star, \boldsymbol{s}^\star)$ of this system by applying variants of Newton's method to the three equalities (3),(4) and (7) and modifying the search directions and step lengths so that the inequalities $(\boldsymbol{x}, \boldsymbol{s}) \geq 0$ are satisfied strictly at every iteration. The equations (3),(4) and (7) are linear or only mildly nonlinear and so are not difficult to solve by themselves. However, the problem becomes much more difficult when we add the nonnegativity requirement $(\boldsymbol{x}, \boldsymbol{s}) \geq 0$, which gives rise to all the complications in the design and analysis of interior-point methods.

To derive primal-dual interior-point methods we restate the optimality conditions in a slightly different form by means of a mapping F from $\mathbb{R}^{2n+m}$ to $\mathbb{R}^{2n+m}$:

$$F(\boldsymbol{x}, \lambda, \boldsymbol{s}) = 0, \tag{10}$$
$$(\boldsymbol{x}, \boldsymbol{s}) \geq 0, \tag{11}$$

Denote
$$X = \operatorname{diag}(x_1, x_2, \cdots, x_n), \quad S = \operatorname{diag}(s_1, s_2, \cdots, s_n),$$

and $\boldsymbol{e} = (1, 1, \cdots, 1)^T$. Primal-dual methods generate iterates $(\boldsymbol{x}^k, \lambda^k, \boldsymbol{x}^k)$ that satisfy the bounds (11) strictly, that is, $\boldsymbol{x}^k > 0$ and $\boldsymbol{s}^k > 0$. This property is the origin of the term interior-point. By respecting these bounds, the methods avoid spurious solutions, that is, points that satisfy $F(\boldsymbol{x}, \lambda, \boldsymbol{s}) = 0$ but not $(\boldsymbol{x}, \boldsymbol{s}) \geq 0$. Like most iterative algorithms in optimization, primal-dual interior-point methods have two basic ingredients: a procedure for determining the step and a measure of the desirability of each point in the search space. An important component of the measure of desirability is the average value of the pairwise products $x_i s_i, i = 1, 2, \cdots, n$, which are all positive when $\boldsymbol{x} > 0$ and $\boldsymbol{s} > 0$. This quantity is known as the duality measure and is defined as follows:

$$\mu = \frac{1}{n} \sum_{i=1}^{n} x_i s_i = \frac{\boldsymbol{x}^T \boldsymbol{s}}{n}. \tag{12}$$

The procedure for determining the search direction has its origins in Newton's method for the nonlinear equations (10). Newton's method forms a linear model for F around the current point and obtains the search direction $(\delta \boldsymbol{x}, \delta \lambda, \delta \boldsymbol{s})$ by solving the following system of linear equations:

$$J(\boldsymbol{x}, \lambda, \boldsymbol{s}) \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \lambda \\ \delta \boldsymbol{s} \end{bmatrix} = -F(\boldsymbol{x}, \lambda, \boldsymbol{s}),$$

where J is the Jacobian of F. If we use the notation $\boldsymbol{r}_c$ and $\boldsymbol{r}_b$ for the first two block rows in F, that is,

$$\boldsymbol{r}_b = A\boldsymbol{x} - \boldsymbol{b}, \quad \boldsymbol{r}_c = A^T \lambda + \boldsymbol{s} - \boldsymbol{c}, \tag{13}$$

we can write the Newton equations as follows:

$$\begin{bmatrix} 0 & A^T & I \\ A & 0 & 0 \\ S & 0 & X \end{bmatrix} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \lambda \\ \delta \boldsymbol{s} \end{bmatrix} = \begin{bmatrix} -\boldsymbol{r}_c \\ -\boldsymbol{r}_b \\ -XS\boldsymbol{e} \end{bmatrix}. \tag{14}$$

Usually, a full step along this direction would violate the bound $(\boldsymbol{x}, \boldsymbol{s}) \geq 0$, so we perform a line search along the Newton direction and define the new iterate as

$$(\boldsymbol{x}, \lambda, \boldsymbol{s}) + \alpha(\delta \boldsymbol{x}, \delta \lambda, \delta \boldsymbol{s}),$$

for some line search parameter $\alpha \in (0, 1]$. We often can take only a small step along this direction $\alpha << 1$ before violating the condition $(\boldsymbol{x}, \boldsymbol{s}) > 0$. Hence, the pure Newton direction (14), sometimes known as the affine scaling direction, often does not allow us to make much progress toward a solution. Most primal-dual methods use a less aggressive Newton direction, one that does not aim directly for a solution but rather for a point whose pairwise products $x_i s_i$ are reduced to a lower average value – not all the way to zero. Specifically, we take a Newton step toward the a point for which $x_i s_i = \sigma \mu$, where μ is the current duality measure and $\sigma \in [0, 1]$ is the reduction factor that we wish to achieve in the duality measure on this step. The modified step equation is then

$$\begin{bmatrix} 0 & A^T & I \\ A & 0 & 0 \\ S & 0 & X \end{bmatrix} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \lambda \\ \delta \boldsymbol{s} \end{bmatrix} = \begin{bmatrix} -\boldsymbol{r}_c \\ -\boldsymbol{r}_b \\ -XS\boldsymbol{e} + \sigma \mu \boldsymbol{e} \end{bmatrix}. \tag{15}$$

We call σ the centering parameter, for reasons to be discussed later in this class. When $\sigma > 0$, it usually is possible to take a longer step α along the direction defined by (15) before violating the bounds $(\boldsymbol{x}, \boldsymbol{s}) \geq 0$. At this point, we have specified most of the elements of a path-following primal-dual interior-point method.

The choices of centering parameter σ_k and step length α_k are crucial to the performance of the method. Techniques for controlling these parameters, directly and indirectly, give rise to a wide variety of methods with diverse properties.

3   Convergence of Interior Point Methods

This section works through the proof of a convergence analysis of an interior point method.

The Central Path. The primal-dual feasible set $\mathcal{F}$ and *strictly feasible set* $\mathcal{F}^o$ are defined as follows:

$$\mathcal{F} = \{(\boldsymbol{x}, \lambda, \boldsymbol{s}) | A\boldsymbol{x} = \boldsymbol{b}, A^T\lambda + \boldsymbol{s} = \boldsymbol{c}, (\boldsymbol{x}, \boldsymbol{s}) \geq 0\}, \tag{16}$$

$$\mathcal{F}^o = \{(\boldsymbol{x}, \lambda, \boldsymbol{s}) | A\boldsymbol{x} = \boldsymbol{b}, A^T\lambda + \boldsymbol{s} = \boldsymbol{c}, (\boldsymbol{x}, \boldsymbol{s}) > 0\}. \tag{17}$$

The central path $\mathcal{C}$ is an arc of strictly feasible points that plays a vital role in primal-dual algorithms. It is parametrized by a scalar $\tau > 0$, and each point $(\boldsymbol{x}^\tau, \lambda^\tau, \boldsymbol{s}^\tau) \in \mathcal{C}$ satisfies the following equations:

$$A^T\lambda + \boldsymbol{s} = \boldsymbol{c}, \tag{18}$$

$$A\boldsymbol{x} = \boldsymbol{b}, \tag{19}$$

$$x_i s_i = \tau, \quad i = 1, 2, \cdots, n, \tag{20}$$

$$(\boldsymbol{x}, \boldsymbol{s}) > 0 \tag{21}$$

These conditions differ from the KKT conditions only in the term τ on the right-hand side of (20). Instead of the complementarity condition (20), we require that the pairwise products $x_i s_i$ have the same (positive) value for all indices i. From (18)-(21), we can define the central path as $\mathcal{C} = \{(\boldsymbol{x}^\tau, \lambda^\tau, \boldsymbol{s}^\tau) | \tau > 0\}$. It can be shown that $(\boldsymbol{x}^\tau, \lambda^\tau, \boldsymbol{s}^\tau)$ is defined uniquely for each $\tau > 0$ if and only if $\mathcal{F}^o$ is nonempty. The conditions (18)-(21) are also the optimality conditions for a logarithmic-barrier formulation. By introducing log-barrier terms for the nonnegativity constraints, with barrier parameter $\tau > 0$, we obtain

$$\min \ \boldsymbol{c}^T\boldsymbol{x} - \tau \sum_{i=1}^n \log(x_i), \quad \text{subject to} \quad A\boldsymbol{x} = \boldsymbol{b}. \tag{22}$$

Central Path Neighborhoods and Path-Following Methods. Path-following algorithms explicitly restrict the iterates to a neighborhood of the central path $\mathcal{C}$ and follow $\mathcal{C}$ to a solution of the linear program. By preventing the iterates from coming too close to the boundary of the nonnegative orthant, they ensure that it is possible to take a nontrivial step along each search direction. Moreover, by forcing the duality measure μ_k to zero as $k \to \infty$, we ensure that the iterates $(\boldsymbol{x}^k, \lambda^k, \boldsymbol{s}^k)$ come closer and closer to satisfying the KKT conditions.

The two most interesting neighborhoods of $\mathcal{C}$ are

$$\mathcal{N}_2(\theta) = \{(\boldsymbol{x}, \lambda, \boldsymbol{s}) \in \mathcal{F}^o | \|XS\boldsymbol{e} - \mu\boldsymbol{e}\| \leq \theta\mu\}, \tag{23}$$

for some $\theta \in [0, 1)$, and

$$\mathcal{N}_{-\infty}(\gamma) = \{(\boldsymbol{x}, \lambda, \boldsymbol{s}) \in \mathcal{F}^o | x_i s_i \geq \gamma\mu, \quad i = 1, 2, \cdots, n\}, \tag{24}$$

for some $\gamma \in (0, 1]$. (Typical values of the parameters are $\theta = 0.5$ and $\gamma = 10^{-3}$.) If a point lies in $\mathcal{N}_{-\infty}(\gamma)$, each pairwise product $x_i s_i$ must be at least some small multiple γ of their average value μ. This requirement is actually quite modest, and we can make $\mathcal{N}_{-\infty}(\gamma)$ encompass most of the feasible region $\mathcal{F}$ by choosing γ close to zero. The $\mathcal{N}_2(\theta)$ neighborhood is more restrictive, since certain points in $\mathcal{F}^o$ do not belong to $\mathcal{N}_2(\theta)$ no matter how close θ is chosen to its upper bound of 1. By keeping all iterates inside one or other of these neighborhoods, path-following methods reduce all the pairwise products $x_i s_i$ to zero at more or less the same rate.

Long-Step Path-Following. The pesudo-code of the algorithm we want to discuss is given below:

- Given $\gamma, \sigma_{\min}, \sigma_{\max}$ with $\gamma \in (0, 1), 0 < \sigma_{\min} \leq \sigma_{\max} < 1$, and $(\boldsymbol{x}^0, \lambda^0, \boldsymbol{s}^0) \in \mathcal{N}_{-\infty}(\gamma)$;

- **for** $k = 0, 1, 2, \ldots$,

-      Choose $\sigma_k \in [\sigma_{\min}, \sigma_{\max}]$;

- Solve the following linear system to obtain $(\Delta \boldsymbol{x}^k, \Delta \lambda^k, \Delta \boldsymbol{s}^k)$:

$$\begin{bmatrix} 0 & A^T & I \\ A & 0 & 0 \\ S^k & 0 & X^k \end{bmatrix} \cdot \begin{bmatrix} \Delta \boldsymbol{x}^k \\ \Delta \lambda^k \\ \Delta \boldsymbol{s}^k \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -X^k S^k \boldsymbol{e} + \sigma \mu^k \boldsymbol{e} \end{bmatrix}. \tag{25}$$

- Choose α_k as the largest value of α in $[0, 1]$ such that

$$(\boldsymbol{x}^k(\alpha), \lambda^k(\alpha), \boldsymbol{s}^k(\alpha)) \in \mathcal{N}_{-\infty}(\gamma)$$

- Set $(\boldsymbol{x}^{k+1}, \lambda^{k+1}, \boldsymbol{s}^{k+1}) = (\boldsymbol{x}^k(\alpha^k), \lambda^k(\alpha^k), \boldsymbol{s}^k(\alpha^k))$.

- **end(for)**

Proof Architecture.

Lemma 3.1. *Let $\boldsymbol{u}$ and $\boldsymbol{v}$ be any two vectors in $\mathbb{R}^n$ with $\boldsymbol{u}^T \boldsymbol{v} \geq 0$. Then*

$$\|UV\boldsymbol{e}\| \leq 2^{-\frac{3}{2}} \|\boldsymbol{u} + \boldsymbol{v}\|_2^2,$$

where

$$U = \mathrm{diag}(u_1, \cdots, u_n), \quad V = \mathrm{diag}(v_1, \cdots, v_n).$$

Lemma 3.2. *If $(\boldsymbol{x}, \lambda, \boldsymbol{s}) \in \mathcal{N}_{-\infty}(\gamma)$, then*

$$\|\Delta X \Delta S \boldsymbol{e}\| \leq 2^{-\frac{3}{2}} (1 + \frac{1}{\gamma}) n \mu.$$

Theorem 3.1. *Given the parameters $\gamma, \sigma_{\min}, \sigma_{\max}$, there is a constant δ independent of n such that*

$$\mu_{k+1} \leq (1 - \frac{\delta}{n}) \mu_k,$$

for all $k \geq 0$.

Theorem 3.2. *Given $\epsilon \in (0, 1)$ and $\gamma \in (0, 1)$, suppose the starting point satisfies $(\boldsymbol{x}^0, \lambda^0, \boldsymbol{s}^0) \in \mathcal{N}_{-\infty}(\gamma)$. Then there is an index K with $K = O(n \log(1/\epsilon))$ such that*

$$\mu_k \leq \epsilon \mu_0, \qquad \text{for all } k \geq K.$$