Linear Programming And Its Interior Point Methods

Lalith Durbhakula

November 2025

Introduction

A linear programming problem (LP) is of the form

$$\max \mathbf{c}^{\mathbf{T}}\mathbf{x}$$

$$\text{subject to } \mathbf{Ax} \leq \mathbf{b},\ \mathbf{x} \geq \mathbf{0}$$

where

$$A = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \vdots & a_{mn} \end{bmatrix}$$

is a coefficient matrix, and

$$c = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix}, \qquad x = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}.$$

are vector matrices. The quantity we are trying to maximize therefore becomes

$$\mathbf{c}^{\mathbf{T}}\mathbf{x} = \sum_{i=1}^{n} c_i x_i$$

In the real world, LPs are most commonly used in transportation systems or profit maximization scenarios. For instance, suppose a company sells two goods X and Y with respective counts x_1 and x_2. Then, the profit-maximizing LP can be represented by

$$\max 4x_1 + 3x_2$$

$$\text{subject to } x_1 + x_2 \leq 48, x_1 \geq 0, x_2 \geq 0$$

When all constraints are mapped in $\mathbb{R}^n$, the Fundamental Theorem of Linear Programming states that if an optimal solution exists, then at least one optimal solution occurs at a vertex of the feasible region. This motivates the simplex method, which travels from vertex to vertex along edges of the feasible polytope.

However, the simplex method can cycle in degenerate vertices (where more than n constraints intersect). To avoid this and achieve better theoretical performance, we examine interior-point methods which instead navigate through the interior of the feasible region.

In this paper, we introduce linear programming, review the necessary mathematical background, and present the main ideas behind the polynomial-time interior point methods that solve it efficiently.

Preliminaries

It is necessary to list rudimentary definitions before proceeding with the methods for solving a linear program. Using the form presented in the introduction, we label A as the **constraint matrix**, $\mathbf{b}$ as the **offset vector** and $\mathbf{c}$ as the **objective vector**.

Definition 1: A **convex set** is a set S such that for any $x, y \in S$ and $\lambda \in [0,1]$, we must have

$$\lambda x + (1-\lambda)y \in S \ (1)$$

Here, (1) is the defining characteristic of the line passing through x and y. In an LP, the feasible polytope is assumed to be convex.

Definition 2: A **convex function** is a function f such that for $\lambda \in [0,1]$, f satisfies the inequality

$$f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda)f(y)$$

Definition 3: A **concave function** is a function f such that for $\lambda \in [0,1]$, f satisfies

$$f(\lambda x + (1-\lambda)y) \geq \lambda f(x) + (1-\lambda)f(y)$$

We are now ready to prove the Fundamental Theorem of Linear Programming.

Theorem: The **fundamental theorem of linear programming** states that given a LP, the global optimization of the objective function exists at a vertex of the feasible convex polytope.

Proof: We note that $f(x) = \mathbf{c}^\mathbf{T} x$ satisfies

$$f(\lambda x + (1-\lambda)y) = \mathbf{c}^\mathbf{T}(\lambda x + (1-\lambda)y)$$
$$= \mathbf{c}^\mathbf{T}(\lambda x) + \mathbf{c}^\mathbf{T}((1-\lambda)y) \text{(via linearity of dot product)}$$
$$= \lambda \mathbf{c}^\mathbf{T} x + (1-\lambda)\mathbf{c}^\mathbf{T} y$$
$$= \lambda f(x) + (1-\lambda)f(y)$$

Hence, the equality case is achieved, implying that the objective function f is convex and concave. By Bauer's maximum principle, given a concave function on a convex set, the minimum value of the function occurs at a vertex (extreme point). Similarly, given a convex function on a convex set, the maximum value of the function occurs at a vertex (extreme point). Therefore, since f is convex and concave, its global optimum (maximization or minimization) exists at a vertex of the feasible convex polytope.

Primal and Dual Problems

A primal LP problem is written in the form

$$\max \mathbf{c^T x}$$

$$\text{subject to } \mathbf{Ax} \leq \mathbf{b},\ \mathbf{x} \geq \mathbf{0}$$

Its corresponding dual problem can be written as

$$\min \mathbf{b^T y}$$

$$\text{subject to } \mathbf{Ay} \geq \mathbf{c},\ \mathbf{y} \geq 0$$

To obtain the dual, consider expanding the primal problem:

$$\max \sum_{i=1}^{n} c_i x_i$$

$$\text{subject to}$$

$$a_{11}x_1 + \cdots + a_{1n}x_1 \leq b_1$$

$$a_{21}x_1 + \cdots + a_{2n}x_1 \leq b_2$$

$$\vdots$$

$$a_{m1}x_1 + \cdots + a_{mn}x_1 \leq b_m$$

$$x_1 \geq 0$$

$$x_2 \geq 0$$

$$\vdots$$

$$x_n \geq 0$$

We can add all objective constraints and scale using y to obtain

$$y_1 \cdot (a_{11}x_1 + \cdots + a_{1n}x_n)$$

$$+\ y_2 \cdot (a_{21}x_1 + \cdots + a_{2n}x_n)$$

$$+\cdots$$

$$+\ y_m \cdot (a_{m1}x_1 + \cdots + a_{mn}x_n)$$

$$\leq y_1 b_1 + y_2 b_2 + \cdots + y_m b_m$$

Now, observe that we can factor $x_1, x_2, \ldots, x_n$ from the left-hand side:

$$\sum_{j=1}^{n} \left(\sum_{i=1}^{m} y_i a_{ij} \right) x_j \leq \sum_{i=1}^{m} y_i b_i.$$

For this inequality to hold for all feasible $x \geq 0$, it must be that each coefficient of x_j on the left is at least c_j, otherwise, we could pick a large x_j to violate the inequality. In matrix notation, this is exactly

$$\mathbf{A^T y} \geq \mathbf{c}.$$

This is the key observation: the dual problem arises naturally by asking, "What combination of constraints gives a guaranteed upper bound on the primal objective?"

Therefore, the dual problem becomes

$$\min \mathbf{b^T y} \quad \text{subject to } \mathbf{A^T y} \geq \mathbf{c}, \quad y \geq 0.$$

This derivation highlights why duality is so important: any feasible dual solution gives an upper bound on the primal, and if we find a pair of primal and dual solutions with matching objective values, we know we have reached optimality. Later, when we discuss interior point methods, we will see that tracking both primal and dual solutions together is the heart of the algorithm, as it allows us to measure how close we are to optimality via the duality gap.

Weak and Strong Duality

Weak and strong duality formalize the relationship between primal and dual problems and provide the theoretical backbone for interior point methods.

Weak Duality: Suppose $\mathbf{x}$ is feasible for the primal and $\mathbf{y}$ is feasible for the dual. Then the primal objective is always less than or equal to the dual objective:

$$\mathbf{c^T x} \leq \mathbf{b^T y}.$$

Proof: Let $\mathbf{x}$ satisfy $\mathbf{Ax} \leq \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$, and let $\mathbf{y}$ satisfy $\mathbf{A^T y} \geq \mathbf{c}$ and $\mathbf{y} \geq \mathbf{0}$. Consider

$$c^T x \leq (\mathbf{A^T y})^T x = y^T(\mathbf{Ax}) \leq y^T b.$$

The first inequality uses dual feasibility: $A^T y \geq c$ implies $(A^T y)^T x \geq c^T x$ for $x \geq 0$. The second inequality uses primal feasibility and $y \geq 0$: each term $y_i(Ax)_i \leq y_i b_i$. Combining these gives $c^T x \leq b^T y$, which is exactly the weak duality statement.

Weak duality is powerful because it immediately gives us a way to verify bounds. Even if we do not know the optimum, weak duality provides a guaranteed inequality.

Strong Duality: Weak duality only gives an inequality. Strong duality tells us that under reasonable assumptions (like strict feasibility), the inequality becomes equality at optimality:

$$\mathbf{c^T x}^* = \mathbf{b^T y}^*.$$

Proof Sketch: For linear programs, this follows from complementary slackness conditions:

$$\begin{aligned}
&\mathbf{Ax}^* \leq \mathbf{b}, \quad x^* \geq 0, \\
&\mathbf{A^T y}^* \geq \mathbf{c}, \quad y^* \geq 0, \\
&y_i^*(b_i - (Ax^*)_i) = 0, \quad x_j^*((A^T y^*)_j - c_j) = 0.
\end{aligned}$$

These conditions say that either a constraint is tight, or its corresponding dual multiplier is zero. When these hold, the duality gap

$$b^T y^* - c^T x^*$$

is zero, proving equality of primal and dual objectives.

Strong duality is critical for interior point methods because it allows us to define the duality gap as a measure of convergence. As the method iterates, we decrease this gap, and once it is sufficiently small, we know we are essentially at the optimum.

Newton's Method for the Barrier Problem

Interior point methods differ from simplex in that they do not move along edges of the polytope but rather traverse its interior. To stay strictly feasible, we introduce a logarithmic barrier for the inequality constraints:

$$\phi(x) = -\sum_{i=1}^{m} \ln(b_i - (Ax)_i) - \sum_{j=1}^{n} \ln x_j.$$

Motivation for the logarithmic barrier: The logarithmic function has the property that it approaches $-\infty$ as its argument approaches zero from the right. By taking the negative log of the slack $b_i - (Ax)_i$ and of each variable x_j, we ensure that as x approaches the boundary of the feasible region, $\phi(x) \to +\infty$. This effectively penalizes solutions near the boundary, keeping iterates safely inside the interior. In addition, ϕ is smooth and strictly convex, which makes it compatible with Newton's method.

The barrier-augmented objective is

$$f_\mu(x) = c^T x + \mu\phi(x),$$

where $\mu > 0$ controls how strongly we enforce staying away from the boundary. As $\mu \to 0$, the barrier becomes negligible and $f_\mu(x)$ approaches the original linear objective.

Newton's method is applied to this smooth, strictly concave (or convex for minimization) objective. At each iteration, we compute the gradient and Hessian:

$$\nabla f_\mu(x) = c - \sum_{i=1}^{m} \frac{a_i}{b_i - a_i^T x} - X^{-1}\mathbf{1}, \quad \nabla^2 f_\mu(x) = \sum_{i=1}^{m} \frac{a_i a_i^T}{(b_i - a_i^T x)^2} + X^{-2}.$$

Here, a_i is the i-th row of A, and X is the diagonal matrix with x on its diagonal.

The Newton step Δx solves

$$\nabla^2 f_\mu(x)\Delta x = -\nabla f_\mu(x),$$

and we update $x \leftarrow x + \alpha\Delta x$, where α is chosen to maintain strict feasibility. Because f_μ is self-concordant, Newton's method achieves quadratic convergence near the optimum. This combination of barrier and Newton step is the core mechanism of interior point methods.

Central Path

The central path is defined as the set of solutions to the barrier problem as μ decreases:

$$x(\mu) = \arg\max\{c^T x + \mu\phi(x)\}.$$

Some important properties:

- The iterates stay strictly feasible: $Ax(\mu) < b$, $x(\mu) > 0$ for all $\mu > 0$.
- As $\mu \to 0$, $x(\mu)$ approaches an optimal solution of the original LP.

Intuitively, the barrier prevents the iterates from touching the boundary too early. By gradually decreasing μ and following Newton directions, the method traces a smooth path from the interior toward the optimum. Each step is designed to reduce the duality gap and maintain strict feasibility.

The central path also provides a geometric perspective: it is a trajectory inside the feasible polytope that guides us directly to the optimum without getting stuck at vertices. This contrasts with the simplex method, which jumps along edges. The central path allows for a polynomial-time analysis and explains why interior point methods are globally convergent.

Polynomial-Time Guarantees and Open Problems

One of the major theoretical breakthroughs of interior point methods is the polynomial-time guarantee. Karmarkar's algorithm and subsequent variants solve a linear program with n variables and m constraints in a number of iterations polynomial in n, m, and the input bit-length L. More precisely, the number of Newton iterations needed to reduce the duality gap below ϵ is $O(\sqrt{n}\log(1/\epsilon))$, and each iteration involves solving a linear system based on the Hessian of the barrier-augmented objective.

This is significant because it contrasts with the simplex method, which can take exponential time in the worst case. The polynomial-time result shows that, in theory, linear programming can always be solved efficiently regardless of the problem size or degeneracy.

Despite this success, there remain important open problems:

- **Strongly polynomial LP:** Can an algorithm exist whose runtime depends only on n and m, independent of input size? This is still unknown.
- **Numerical stability and efficiency:** For very large-scale LPs, designing interior point methods that are both fast and numerically robust remains challenging.
- **Extensions to other problems:** Applying interior point methods to quadratic, semidefinite, and general convex optimization remains an active research area, with practical and theoretical questions about convergence and complexity.

In summary, interior point methods give us both practical efficiency and strong theoretical guarantees. By following the central path with barrier functions and Newton steps, tracking the duality gap, and exploiting convexity, these methods provide a powerful alternative to simplex, with room for ongoing improvements and deeper theoretical understanding.

References

1. Komei Fukuda, "Linear Programming" (Lecture Notes, Eidgenössische Technische Hochschule [ETH] Zürich), accessed December 8, 2025, https://people.inf.ethz.ch/fukudak/lect/opt2011/aopt11note4.pdf.

2. Michel X. Goemans, "18.310 Lecture 3: Simplex Method and Duality" (Course Notes, Massachusetts Institute of Technology, Spring 2015), accessed December 8, 2025, https://math.mit.edu/ goemans/18310S15/lpnotes310.pdf.

3. Jon Kleinberg and Éva Tardos, "Chapter 4: Network Flow: The Maximum-Flow Problem" (Course Material, Massachusetts Institute of Technology), accessed December 8, 2025, http://web.mit.edu/15.053/www/AMP-Chapter-04.pdf.

4. David Mount, "Linear Programming" (Course Notes, Department of Computer Science, University of Illinois at Urbana-Champaign, 2015), accessed December 8, 2025, https://jeffe.cs.illinois.edu/teaching/algorithms/notes/H-lp.pdf.