LINEAR PROGRAMS

CSCI 4113/6101
INSTRUCTOR: NORBERT ZEH
SEPTEMBER 9, 2025

Most of this course is concerned with developing algorithms to solve **optimization problems**, and virtually all optimization problems we study in this course can be expressed as **linear programs** (LPs). This topic introduces you to optimization problems and linear programs.

We are interested in linear programs primarily as a language to formally define optimization problems. Apart from providing an unambiguous definition of the problem we aim to solve, the linear program is often the starting point for efficient algorithms to solve the problem it defines. In an upcoming topic, we will discuss the Simplex Algorithm as a general algorithm to solve arbitrary linear programs. The Simplex Algorithm can be viewed as a numerical algorithm that has no understanding of the problem described by the linear program it is given as input. More often, we are interested in purely combinatorial, problem-specific algorithms. Interestingly, the LP formulation—especially via LP duality, a concept discussed in an upcoming topic—often reveals valuable information about the structure of the problem, which can be exploited to obtain efficient combinatorial algorithms for the problems we are interested in.

1 OPTIMIZATION PROBLEMS

An optimization problem has possibly many "correct" solutions. Our goal is to pick the best solution among them according to some quality measure associated with each solution. More precisely, I should say "*a* best solution" because, in general, there can be more than one solution that are all equally good and better than all other solutions.

Consider the problem of sorting a list of numbers in increasing order. There is only one correct output for any given input, so this is not an optimization problem. On the other hand, *designing* a sorting algorithm can be viewed as an optimization problem: There are many correct sorting algorithms, but our goal usually is to choose the algorithm with the lowest running time. Our quality measure in this case is the running time of the algorithm, and our goal is to find a solution—that is, a sorting algorithm—that minimizes this quality measure.

Formally, we describe an optimization problem as the problem of assigning real values to a set of variables $x_1, \ldots, x_n$ such that some **objective function** $f(x_1, \ldots, x_n)$ is maximized or minimized, subject to some **constraints** that the values assigned to these variables must satisfy. If our goal is to maximize $f(x_1, \ldots, x_n)$, then we call the problem a **maximization problem**; if our goal is to minimize $f(x_1, \ldots, x_n)$, a **minimization problem**. We usually write the list of variables more compactly as a (column) vector $x = (x_1, \ldots, x_n)^T$.[1] We will then also treat f simply as a function of this vector, $f(x)$.

The next bit of terminology is something to get used to, so pay attention: A **solution** of the optimization problem is simply an assignment of values to the variables $x_1, \ldots, x_n$. It *does not* need to

[1] x^T is the transpose of the vector x.

minimize or maximize $f(x)$; it doesn't even need to satisfy the constraints imposed by the problem. We refer to such a solution as a vector $\hat{x} = (\hat{x}_1, \ldots, \hat{x}_n)$, where $\hat{x}_i$ is the value assigned to variable x_i, for all $i \in [n]$.[2] We call $\hat{x}$ a **feasible solution** if it satisfies all the constraints imposed by the problem. A feasible solution is an **optimal solution** if at maximizes or minimizes $f(\hat{x})$, according to whether the problem is a maximization or minimization problem.

Here is a really simple example that illustrates all of these concepts. Every rectangular prism in 3-d can be described by its three side lengths x, y, z. These are our variables. Now let us assume that we want a prism whose total side length $x + y + z$ is at most 15. In addition side lengths obviously cannot be negative. We can express these conditions using four constraints:

$$\begin{aligned} x + y + z &\leq 15 \\ x &\geq 0 \\ y &\geq 0 \\ z &\geq 0 \end{aligned}$$

Our objective is to find a rectangular prism that satisfies these four constraints and has maximum volume. Thus, our objective function is

$$f(x, y, z) = x \cdot y \cdot z.$$

Since our goal is to maximize this objective function, the problem is a maximization problem. The following three are all solutions of this problem:

$$\begin{aligned} (\hat{x}, \hat{y}, \hat{z})^T &= (15, -3, 20)^T \\ (\tilde{x}, \tilde{y}, \tilde{z})^T &= (3, 4, 8)^T \\ (\dot{x}, \dot{y}, \dot{z})^T &= (5, 5, 5)^T \end{aligned}$$

$(\hat{x}, \hat{y}, \hat{z})^T$ is not a feasible solution because $\hat{y} < 0$ and $\hat{x} + \hat{y} + \hat{z} > 15$. $(\tilde{x}, \tilde{y}, \tilde{z})^T$ and $(\dot{x}, \dot{y}, \dot{z})^T$ are both feasible solutions. Since $3 \cdot 4 \cdot 8 = 96 < 125 = 5 \cdot 5 \cdot 5$, $(\tilde{x}, \tilde{y}, \tilde{z})^T$ is not an optimal solution. Using a bit of calculus, we would be able to verify that $(\dot{x}, \dot{y}, \dot{z})^T$ *is* an optimal solution. This just captures the well-known fact that given fixed bound on its total side length, the rectangular prism of maximum volume is a cube.

2 Linear Programs

A **linear program** (LP) is an optimization problem in which the objective function is a linear combination of the variables, and all constraints are linear inequalities.

A **linear combination** of variables $x_1, \ldots, x_n$ is an expression of the form

$$\sum_{i=1}^{n} a_i x_i,$$

where $a_1, \ldots, a_n$ are arbitrary constants. If you remember linear algebra, then this is nothing but the

[2]Throughout this course, we will use the notation $[n] = \{1, \ldots, n\}$ and $[n]_0 = \{0, \ldots, n\}$.

inner product

$$a^T x,$$

where $a = (a_1, \dots, a_n)^T$.

A **linear inequality** is an expression of the form

$$\begin{aligned} f(x_1, \dots, x_n) &= c, \\ f(x_1, \dots, x_n) &\le c \text{ or} \\ f(x_1, \dots, x_n) &\ge c, \end{aligned}$$

where $f(x_1, \dots, x_n)$ is a linear combination of $x_1, \dots, x_n$ and c is a constant. Note that we disallow strict inequalities. This is because, in general, we can guarantee that there exists an optimal solution only if the space of feasible solutions is a closet subset of $\mathbb{R}^n$.

Our cube finding problem from the previous section *is not* an LP. Its constraints are all linear inequalities, but its objective function $x \cdot y \cdot z$ is not a linear combination of the three variables x, y, and z.

The following *is* a linear program (with only two variables x and y), is you can easily verify that the objective function is a linear combination of x and y, and all constraints are linear constraints:

$$\begin{aligned} &\text{Maximize } x + 2y \\ \text{s.t. (subject to)} \quad & x \ge 0 \\ & y \ge 0 \\ & 3x + y \le 26 \\ & 4x - y \le 31 \\ & 2y - x \le 4 \end{aligned} \tag{1}$$

3 The Geometry of Linear Programs

Every solution $\hat{x} = (\hat{x}_1, \dots, \hat{x}_n)^T$ of an LP with n variables can be interpreted as a point in $\mathbb{R}^n$, and every point in $\mathbb{R}^n$ can be interpreted as a solution of such an LP. This raises the question of which points in $\mathbb{R}^n$ are feasible solutions, and which points are optimal solutions of the LP. This geometric view of feasible and optimal solutions is at the heart of the Simplex Algorithm and of virtually every general algorithm for solving linear programs.

Let us start with a single constraint. If it is is an equality constraint $a^T x = b$, and we assume that $a \ne 0$,[3] then the points that satisfy this constraint form an $(n-1)$-dimensional subspace of $\mathbb{R}^n$, also

[3]Throughout this course, we will write 0 and 1 to denote vectors $0 = (0, \dots, 0)^T$ and $1 = (1, \dots, 1)^T$ of the appropriate dimensions. We will also write

$$1 = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix},$$

that is, 1 may also denote the identity matrix. Whether 0 or 1 refers to a single value, a vector or a matrix will be clear from the context.

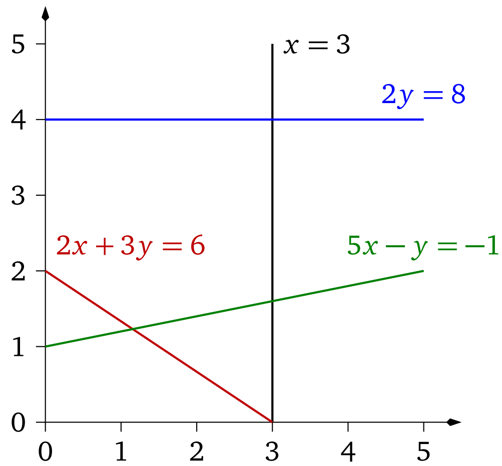

Figure 1: The three lines defined by the constraints $x = 3$ (black), $2y = 8$ (blue), $2x + 3y = 6$ (red), and $5x - y = -1$ (green). Note in particular that the last two eqations can be rewritten in a form that may feel more familiar: $y = -x/3 + 2$ and $y - x/5 + 1$.

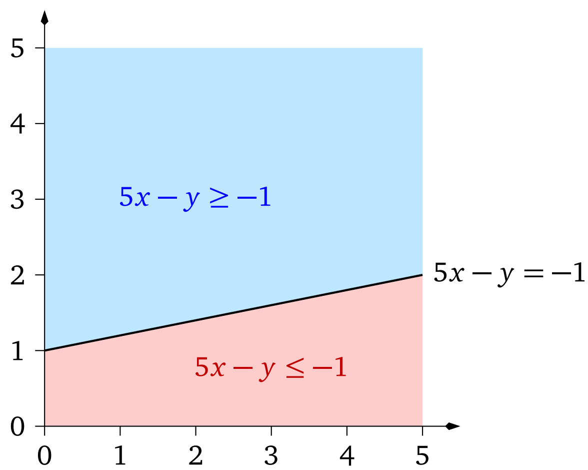

Figure 2: The line $5x - y = -1$ slices the plane into two halves composed of all points satisfying $5x - y \leq -1$ and $5x - y \geq -1$, respectively.

called a **hyperplane**. Fig. 1 shows a few examples in $n = 2$ dimensions, where a hyperplanet is simply a line.

We usually do not care about the case when $a = 0$, as this means that the values assigned to the variables have no impact on whether the constraint is satisfied or not. If the constraint is $0^T x = 0$, then all solutions satisfy this constraint, that is, the constraint imposes no restrictions on the set of feasible solutions and can be omitted. If the constraint is $0^T x = c$, for some $c \neq 0$, then no solution satisfies this constraint, that is, there is no feasible solution. Therefore, both types of constraints are quite useless.

If the constraint is an inequality, then the solutions that satisfy this constraint form an n-dimensional **halfspace**, that is, the set of all points on one side of a hyperplane. In $n = 2$ dimensions, this is easy to visualize once again. A hyperplane is a line in this case. This line slices $\mathbb{R}^2$ into two halves, the two halfspaces on either side of the line. Fig. 2 illustrates this.

Now, a linear program usually has more than one constraint. Assuming these constraints are all

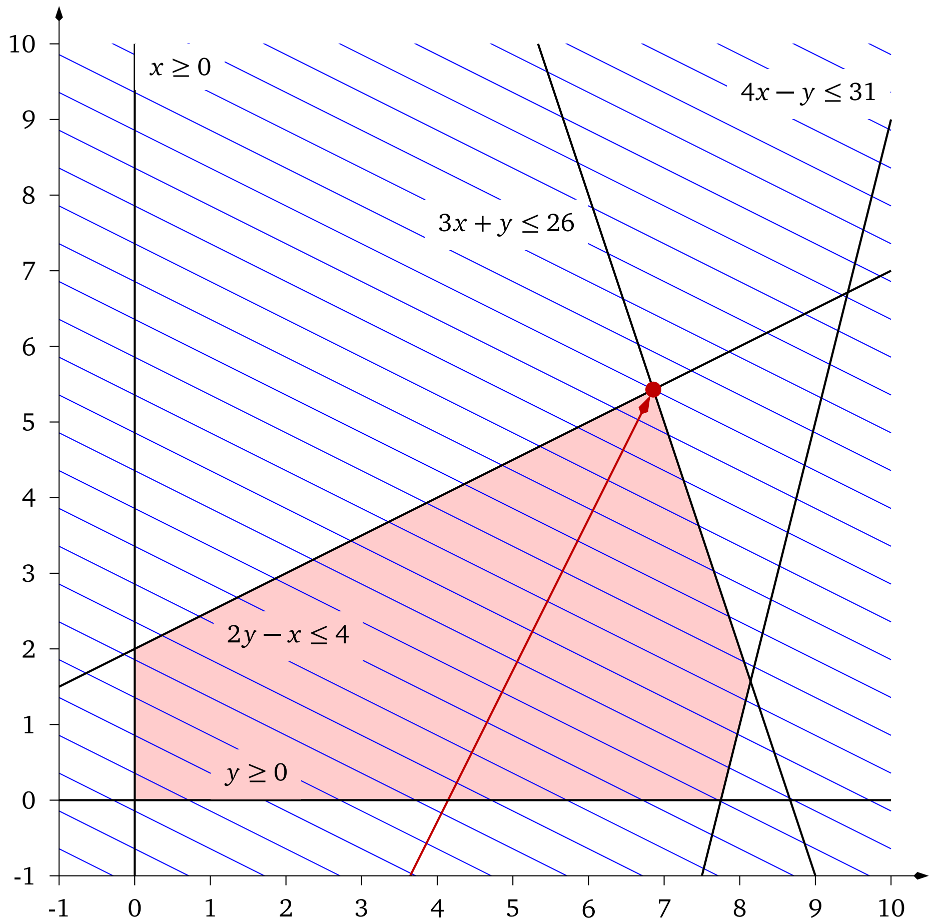

Figure 3: The feasible region of the LP (1) and the optimal solution (red dot). Each of the blue lines is defined by the equation $x + 2y = c$, for some constant c, that is, all of the points on such a line achieve the same objective function value c. As we move these lines upwards, the objective function value increases. As long as a line $x + 2y = c$ has an intersection with the feasible region of the LP, there exists a feasible solution with objective function value c. In particular, if the line intersects the *interior* of the feasible region, there exists a feasible solution with obljective function value greater than c, which lies on a blue line just slightly above the current line. The optimal solution is obtained as soon as the current blue line intersects only the boundary of the feasible region. In this example, this intersection is a single point, that is, this LP has exactly one optimal solution. This is not always the case.

inequality constraints, the points that satisfy each constraint form a halfspace, as just observed. The points that satisfy *all* constraints are exactly the intersection of these halfspaces. The intersection of halfspaces is easily seen to form a convex polytope in $\mathbb{R}^n$. We call this region the **feasible region** of the LP, because a point in $\mathbb{R}^n$ is a feasible solution of the LP if and only if it lies inside this region. Fig. 3 shows the feasible region of the LP (1).

The optimal solution of this LP is the red point $\left(\frac{48}{7}, \frac{38}{7}\right)$ in Fig. 3 with objective function value $\frac{48}{7} + 2 \cdot \frac{38}{7} = \frac{124}{7}$.

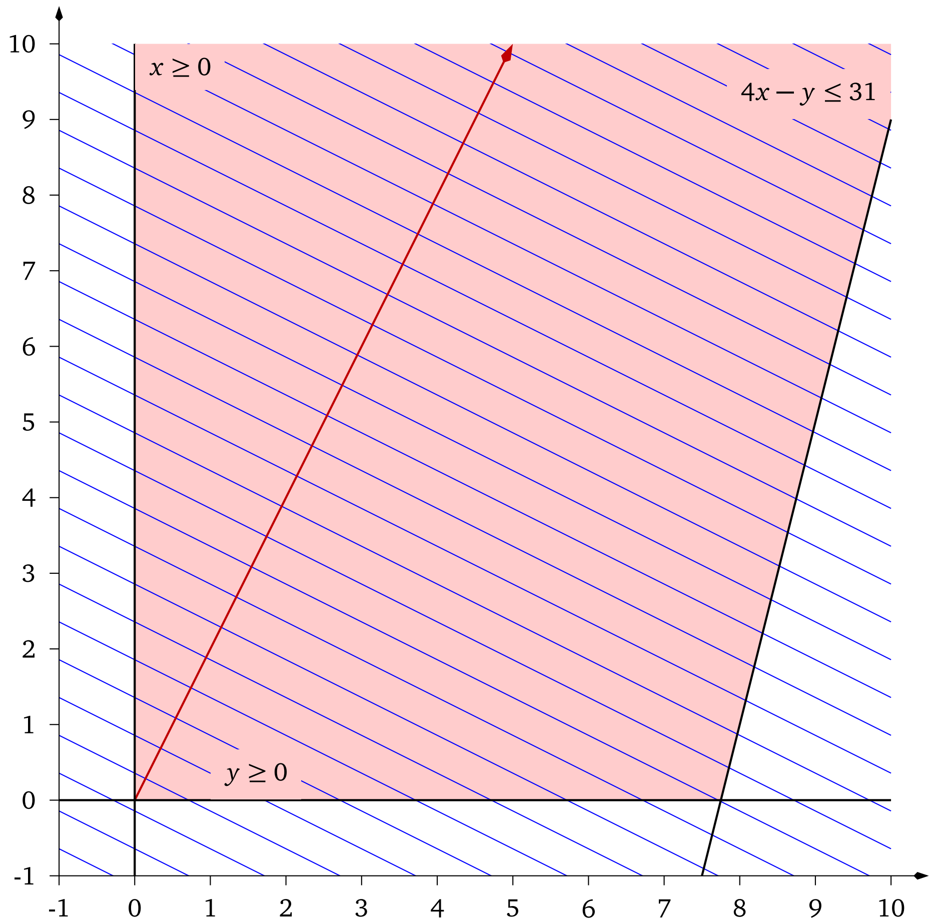

Figure 4: The feasible region of the LP (2). Note that this feasible region is unbounded, as it extends infinitely in the $+y$-direction. Since the objective function value increases as we move in the direction of the red arrow, this means that for every feasible solution, there exists a better feasible solution that is slightly above it. Therefore, this LP does not have an optimal solution.

4 Existence and Properties of an Optimal Solution

Since we are interested in finding optimal solutions of LPs, it is helpful to understand when such solutions exist and, if they exist, what properties characterize them.

First, if the feasible region of the LP is empty, then by definition, there exists no feasible solution and, therefore, no optimal solution. We call the LP **infeasible** in this case. A trivial example is an LP with two constraints $x \leq 1$ and $x \geq 2$. Clearly, there is no solution that satisfies both constraints. Geometrically, these two constraints define two disjoint halfspaces, so their intersection is empty. In general, infeasible LPs arise from more complex interactions between constraints. It is often the case that any two of the constraints have feasible regions with a non-empty intersection, but once we intersect the feasible regions of *all* constraints, this intersection is empty.

When the LP is infeasible, there is no optimal solution because of a lack of feasible solutions. We may also be unable to find an optimal solution because of an overabundance of feasible solutions. If we drop the two constraints $3x + y \leq 26$ and $2y - x \leq 4$ from the LP (1), then we obtain the following LP:

$$\begin{aligned} & \text{Maximize } x + 2y \\ \text{s.t. (subject to)} \quad & x \geq 0 \\ & y \geq 0 \\ & 4x - y \leq 31 \end{aligned} \tag{2}$$

Its feasible region is shown in Fig. 4. This feasible region extends infinitely in the direction that increases the objective function. Thus, for every feasible solution, we can find a better feasible solution; there is no upper bound on the objective function values of feasible solutions. We call such an LP **unbounded** because it has no bound on the achievable objective function value.

While every unbounded LP has an unbounded feasible region, the converse is not true. For example, if we dropped the two constraints $x \geq 0$ and $y \geq 0$ from the LP (1), then the feasible region of the LP would extend infinitely in the $-x$ and $-y$-directions. Thus, the feasible region is unbounded. However, you can verify that the red point in Fig. 3 remains the optimal solution of the LP because the constraints we did not eliminate still limit how far we can move in the direction that improves the objective function value while keeping the solution feasible.

Let us assume then that the LP we want to solve is feasible and bounded. Then we have

Proposition 1. *All optimal solutions of a feasible and bounded LP with a non-constant objective function are points on the boundary of the feasible region of the LP.*

Proof. Consider a feasible solution $\hat{x}$ that is not on the boundary of the feasible region R of the LP. Then there exists a small radius $\varepsilon > 0$ such that every solution $\tilde{x}$ at distance at most ε from $\hat{x}$ is also feasible. Formally, any such solution satisfies

$$\|\hat{x} - \tilde{x}\|_2 = \sqrt{\sum_{i=1}^{n} (\hat{x}_i - \tilde{x}_i)^2} \leq \varepsilon.$$

Let the objective function of the LP be $c^T x$. Since the objective function is non-constant, there exists a variable x_i whose objective function coefficient c_i is non-zero. Assume w.l.o.g. that $c_i > 0$ and that our goal is to maximize the objective function $c^T x$. Then consider the solution $\tilde{x}$ defined as

$$\tilde{x}_j = \begin{cases} \hat{x}_j + \varepsilon & \text{if } j = i \\ \hat{x}_j & \text{otherwise.} \end{cases}$$

This is a feasible solution because

$$\|\hat{x} - \tilde{x}\|_2 = \sqrt{\sum_{j=1}^{n} (\hat{x}_j - \tilde{x}_j)^2} = \sqrt{(\hat{x}_i - \tilde{x}_i)^2} = |\hat{x}_i - \tilde{x}_i| = \varepsilon.$$

Its objective function value is

$$c^T \tilde{x} = c^T \hat{x} + c_i \varepsilon > c^T \hat{x}.$$

This proves that any feasible solution $\hat{x}$ not on the boundary of R is not an optimal solution. Its contrapositive is the lemma. □