

*The greatest flood has the soonest ebb;
the sorest tempest the most sudden calm;
the hottest love the coldest end; and
from the deepest desire oftentimes ensues the deadliest hate.*

— Socrates

*Th' extremes of glory and of shame,
Like east and west, become the same.*

— Samuel Butler, *Hudibras* Part II, Canto I (c. 1670)

*Extremes meet, and there is no better example
than the haughtiness of humility.*

— Ralph Waldo Emerson, "Greatness",
in *Letters and Social Aims* (1876)

*25 Linear Programming

The maximum flow/minimum cut problem is a special case of a very general class of problems called *linear programming*. Many other optimization problems fall into this class, including minimum spanning trees and shortest paths, as well as several common problems in scheduling, logistics, and economics. Linear programming was used implicitly by Fourier in the early 1800s, but it was first formalized and applied to problems in economics in the 1930s by Leonid Kantorovich. Kantorovich's work was hidden behind the Iron Curtain (where it was largely ignored) and therefore unknown in the West. Linear programming was rediscovered and applied to shipping problems in the early 1940s by Tjalling Koopmans. The first complete algorithm to solve linear programming problems, called the *simplex method*, was published by George Dantzig in 1947. Koopmans first proposed the name "linear programming" in a discussion with Dantzig in 1948. Kantorovich and Koopmans shared the 1975 Nobel Prize in Economics "for their contributions to the theory of optimum allocation of resources". Dantzig did not; his work was apparently too pure. Koopmans wrote to Kantorovich suggesting that they refuse the prize in protest of Dantzig's exclusion, but Kantorovich saw the prize as a vindication of his use of mathematics in economics, which his Soviet colleagues had written off as "a means for apologists of capitalism".

A linear programming problem asks for a vector $x \in \mathbb{R}^d$ that maximizes (or equivalently, minimizes) a given linear function, among all vectors x that satisfy a given set of linear inequalities. The general form of a linear programming problem is the following:

$$\begin{aligned} & \text{maximize} && \sum_{j=1}^d c_j x_j \\ & \text{subject to} && \sum_{j=1}^d a_{ij} x_j \leq b_i \quad \text{for each } i = 1 \dots p \\ & && \sum_{j=1}^d a_{ij} x_j = b_i \quad \text{for each } i = p + 1 \dots p + q \\ & && \sum_{j=1}^d a_{ij} x_j \geq b_i \quad \text{for each } i = p + q + 1 \dots n \end{aligned}$$

Here, the input consists of a matrix $A = (a_{ij}) \in \mathbb{R}^{n \times d}$, a column vector $b \in \mathbb{R}^n$, and a row vector $c \in \mathbb{R}^d$. Each coordinate of the vector x is called a *variable*. Each of the linear inequalities is called a *constraint*.

The function $x \mapsto c \cdot x$ is called the *objective function*. I will always use d to denote the number of variables, also known as the *dimension* of the problem. The number of constraints is usually denoted n .

A linear programming problem is said to be in *canonical form*¹ if it has the following structure:

$$\begin{aligned} & \text{maximize } \sum_{j=1}^d c_j x_j \\ & \text{subject to } \sum_{j=1}^d a_{ij} x_j \leq b_i \quad \text{for each } i = 1 \dots n \\ & \quad \quad \quad x_j \geq 0 \quad \text{for each } j = 1 \dots d \end{aligned}$$

We can express this canonical form more compactly as follows. For two vectors $x = (x_1, x_2, \dots, x_d)$ and $y = (y_1, y_2, \dots, y_d)$, the expression $x \geq y$ means that $x_i \geq y_i$ for every index i .

$$\begin{array}{l} \max \quad c \cdot x \\ \text{s.t. } Ax \leq b \\ \quad \quad x \geq 0 \end{array}$$

Any linear programming problem can be converted into canonical form as follows:

- For each variable x_j , add the equality constraint $x_j = x_j^+ - x_j^-$ and the inequalities $x_j^+ \geq 0$ and $x_j^- \geq 0$.
- Replace any equality constraint $\sum_j a_{ij} x_j = b_i$ with two inequality constraints $\sum_j a_{ij} x_j \geq b_i$ and $\sum_j a_{ij} x_j \leq b_i$.
- Replace any upper bound $\sum_j a_{ij} x_j \geq b_i$ with the equivalent lower bound $\sum_j -a_{ij} x_j \leq -b_i$.

This conversion potentially double the number of variables and the number of constraints; fortunately, it is rarely necessary in practice.

Another useful format for linear programming problems is *slack form*², in which every inequality is of the form $x_j \geq 0$:

$$\begin{array}{l} \max \quad c \cdot x \\ \text{s.t. } Ax = b \\ \quad \quad x \geq 0 \end{array}$$

It's fairly easy to convert any linear programming problem into slack form. Slack form is especially useful in executing the simplex algorithm (which we'll see in the next lecture).


25.1 The Geometry of Linear Programming

A point $x \in \mathbb{R}^d$ is *feasible* with respect to some linear programming problem if it satisfies all the linear constraints. The set of all feasible points is called the *feasible region* for that linear program. The feasible region has a particularly nice geometric structure that lends some useful intuition to the linear programming algorithms we'll see later.

¹Confusingly, some authors call this *standard form*.

²Confusingly, some authors call this *standard form*.

Any linear equation in d variables defines a *hyperplane* in \mathbb{R}^d ; think of a line when $d = 2$, or a plane when $d = 3$. This hyperplane divides \mathbb{R}^d into two *halfspaces*; each halfspace is the set of points that satisfy some linear inequality. Thus, the set of feasible points is the intersection of several hyperplanes (one for each equality constraint) and halfspaces (one for each inequality constraint). The intersection of a finite number of hyperplanes and halfspaces is called a *polyhedron*. It's not hard to verify that any halfspace, and therefore any polyhedron, is *convex*—if a polyhedron contains two points x and y , then it contains the entire line segment \overline{xy} .




A two-dimensional polyhedron (white) defined by 10 linear inequalities.

By rotating \mathbb{R}^d (or choosing a coordinate frame) so that the objective function points downward, we can express *any* linear programming problem in the following geometric form:

Find the lowest point in a given polyhedron.


With this geometry in hand, we can easily picture two pathological cases where a given linear programming problem has no solution. The first possibility is that there are no feasible points; in this case the problem is called *infeasible*. For example, the following LP problem is infeasible:

$$\begin{aligned} &\text{maximize } x - y \\ &\text{subject to } 2x + y \leq 1 \\ &\quad \quad \quad x + y \geq 2 \\ &\quad \quad \quad x, y \geq 0 \end{aligned}$$



An infeasible linear programming problem; arrows indicate the constraints.

The second possibility is that there are feasible points at which the objective function is arbitrarily large; in this case, we call the problem *unbounded*. The same polyhedron could be unbounded for some objective functions but not others, or it could be unbounded for every objective function.



A two-dimensional polyhedron (white) that is unbounded downward but bounded upward.

25.2 Example 1: Shortest Paths

We can compute the length of the shortest path from s to t in a weighted directed graph by solving the following very simple linear programming problem.

$$\begin{aligned} &\text{maximize} && d_t \\ &\text{subject to} && d_s = 0 \\ &&& d_v - d_u \leq \ell_{u \rightarrow v} \quad \text{for every edge } u \rightarrow v \end{aligned}$$

Here, $\ell_{u \rightarrow v}$ is the length of the edge $u \rightarrow v$. Each variable d_v represents a tentative shortest-path distance from s to v . The constraints mirror the requirement that every edge in the graph must be relaxed. These relaxation constraints imply that in any feasible solution, d_v is *at most* the shortest path distance from s to v . Thus, somewhat counterintuitively, we are correctly *maximizing* the objective function to compute the *shortest* path! In the optimal solution, the objective function d_t is the actual shortest-path distance from s to t , but for any vertex v that is not on the shortest path from s to t , d_v may be an underestimate of the true distance from s to v . However, we can obtain the true distances from s to every other vertex by modifying the objective function:

$$\begin{aligned} &\text{maximize} && \sum_v d_v \\ &\text{subject to} && d_s = 0 \\ &&& d_v - d_u \leq \ell_{u \rightarrow v} \quad \text{for every edge } u \rightarrow v \end{aligned}$$

There is another formulation of shortest paths as an LP minimization problem using an indicator variable $x_{u \rightarrow v}$ for each edge $u \rightarrow v$.

$$\begin{aligned} &\text{minimize} && \sum_{u \rightarrow v} \ell_{u \rightarrow v} \cdot x_{u \rightarrow v} \\ &\text{subject to} && \sum_u x_{u \rightarrow s} - \sum_w x_{s \rightarrow w} = 1 \\ &&& \sum_u x_{u \rightarrow t} - \sum_w x_{t \rightarrow w} = -1 \\ &&& \sum_u x_{u \rightarrow v} - \sum_w x_{v \rightarrow w} = 0 \quad \text{for every vertex } v \neq s, t \\ &&& x_{u \rightarrow v} \geq 0 \quad \text{for every edge } u \rightarrow v \end{aligned}$$

Intuitively, $x_{u \rightarrow v} = 1$ means $u \rightarrow v$ lies on the shortest path from s to t , and $x_{u \rightarrow v} = 0$ means $u \rightarrow v$ does not lie on this shortest path. The constraints merely state that the path should start at s , end at t , and either pass through or avoid every other vertex v . Any path from s to t —in particular, the shortest path—clearly implies a feasible point for this linear program.

However, there are other feasible solutions, possibly even *optimal* solutions, with non-integral values that do not represent paths. Nevertheless, there is always an optimal solution in which every x_e is either 0 or 1 and the edges e with $x_e = 1$ comprise the shortest path. (This fact is by no means obvious, but a proof is beyond the scope of these notes.) Moreover, in any optimal solution, even if not every x_e is an integer, the objective function gives the shortest path distance!

25.3 Example 2: Maximum Flows and Minimum Cuts

Recall that the input to the maximum (s, t) -flow problem consists of a weighted directed graph $G = (V, E)$, two special vertices s and t , and a function assigning a non-negative *capacity* c_e to each edge e . Our task is to choose the flow f_e across each edge e , as follows:

$$\begin{aligned} &\text{maximize} && \sum_w f_{s \rightarrow w} - \sum_u f_{u \rightarrow s} \\ &\text{subject to} && \sum_w f_{v \rightarrow w} - \sum_u f_{u \rightarrow v} = 0 \quad \text{for every vertex } v \neq s, t \\ &&& f_{u \rightarrow v} \leq c_{u \rightarrow v} \quad \text{for every edge } u \rightarrow v \\ &&& f_{u \rightarrow v} \geq 0 \quad \text{for every edge } u \rightarrow v \end{aligned}$$

Similarly, the minimum cut problem can be formulated using ‘indicator’ variables similarly to the shortest path problem. We have a variable S_v for each vertex v , indicating whether $v \in S$ or $v \in T$, and a variable $X_{u \rightarrow v}$ for each edge $u \rightarrow v$, indicating whether $u \in S$ and $v \in T$, where (S, T) is some (s, t) -cut.³

$$\begin{aligned} &\text{minimize} && \sum_{u \rightarrow v} c_{u \rightarrow v} \cdot X_{u \rightarrow v} \\ &\text{subject to} && X_{u \rightarrow v} + S_v - S_u \geq 0 \quad \text{for every edge } u \rightarrow v \\ &&& X_{u \rightarrow v} \geq 0 \quad \text{for every edge } u \rightarrow v \\ &&& S_s = 1 \\ &&& S_t = 0 \end{aligned}$$

Like the minimization LP for shortest paths, there can be optimal solutions that assign fractional values to the variables. Nevertheless, the minimum value for the objective function is the cost of the minimum cut, and there is an optimal solution for which every variable is either 0 or 1, representing an actual minimum cut. No, this is not obvious; in particular, my claim is not a proof!

25.4 Linear Programming Duality

Each of these pairs of linear programming problems is related by a transformation called *duality*. For any linear programming problem, there is a corresponding dual linear program that can be obtained by a mechanical translation, essentially by swapping the constraints and the variables. The translation is simplest when the LP is in canonical form:

Primal (II)	\iff	Dual (II)
$\begin{aligned} \max & \quad c \cdot x \\ \text{s.t.} & \quad Ax \leq b \\ & \quad x \geq 0 \end{aligned}$		$\begin{aligned} \min & \quad y \cdot b \\ \text{s.t.} & \quad yA \geq c \\ & \quad y \geq 0 \end{aligned}$

³These two linear programs are not quite *syntactic* duals; I've added two redundant variables S_s and S_t to the min-cut program to increase readability.

We can also write the dual linear program in exactly the same canonical form as the primal, by swapping the coefficient vector c and the objective vector b , negating both vectors, and replacing the constraint matrix A with its negative transpose.⁴

Primal (Π) $\begin{aligned} \max \quad & c \cdot x \\ \text{s.t.} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$	\iff	Dual (Π) $\begin{aligned} \max \quad & -b^\top \cdot y^\top \\ \text{s.t.} \quad & -A^\top y^\top \leq -c \\ & y^\top \geq 0 \end{aligned}$
---	--------	--

Written in this form, it should be immediately clear that duality is an *involution*: The dual of the dual linear program Π is identical to the primal linear program Π . The choice of which LP to call the ‘primal’ and which to call the ‘dual’ is totally arbitrary.⁵

The Fundamental Theorem of Linear Programming. *A linear program Π has an optimal solution x^* if and only if the dual linear program Π has an optimal solution y^* such that $c \cdot x^* = y^*Ax^* = y^* \cdot b$.*

The weak form of this theorem is trivial to prove.

Weak Duality Theorem. *If x is a feasible solution for a canonical linear program Π and y is a feasible solution for its dual Π , then $c \cdot x \leq yAx \leq y \cdot b$.*

Proof: Because x is feasible for Π , we have $Ax \leq b$. Since y is positive, we can multiply both sides of the inequality to obtain $yAx \leq y \cdot b$. Conversely, y is feasible for Π and x is positive, so $yAx \geq c \cdot x$. \square

It immediately follows that if $c \cdot x = y \cdot b$, then x and y are optimal solutions to their respective linear programs. This is in fact a fairly common way to prove that we have the optimal value for a linear program.

25.5 Duality Example

Before I prove the stronger duality theorem, let me first provide some intuition about where this duality thing comes from in the first place.⁶ Consider the following linear programming problem:

$$\begin{aligned} \text{maximize} \quad & 4x_1 + x_2 + 3x_3 \\ \text{subject to} \quad & x_1 + 4x_2 \leq 2 \\ & 3x_1 - x_2 + x_3 \leq 4 \\ & x_1, x_2, x_3 \geq 0 \end{aligned}$$

Let σ^* denote the optimum objective value for this LP. The feasible solution $x = (1, 0, 0)$ gives us a lower bound $\sigma^* \geq 4$. A different feasible solution $x = (0, 0, 3)$ gives us a better lower bound $\sigma^* \geq 9$. We could play this game all day, finding different feasible solutions and getting ever larger lower bounds. How do we know when we’re done? Is there a way to prove an *upper* bound on σ^* ?

⁴For the notational purists: In these formulations, x and b are column vectors, and y and c are row vectors. This is a somewhat nonstandard choice. Yes, that means the dot in $c \cdot x$ is redundant. Sue me.

⁵For historical reasons, maximization LPs tend to be called ‘primal’ and minimization LPs tend to be called ‘dual’. This is a pointless religious tradition, nothing more. Duality is a relationship between LP problems, not a type of LP problem.

⁶This example is taken from Robert Vanderbei’s excellent textbook *Linear Programming: Foundations and Extensions* [Springer, 2001], but the idea appears earlier in Jens Clausen’s 1997 paper ‘Teaching Duality in Linear Programming: The Multiplier Approach’.

In fact, there is. Let's multiply each of the constraints in our LP by a new non-negative scalar value y_i :

$$\begin{aligned} & \text{maximize} && 4x_1 + x_2 + 3x_3 \\ & \text{subject to} && y_1(x_1 + 4x_2) \leq 2y_1 \\ & && y_2(3x_1 - x_2 + x_3) \leq 4y_2 \\ & && x_1, x_2, x_3 \geq 0 \end{aligned}$$

Because each y_i is non-negative, we do not reverse any of the inequalities. Any feasible solution (x_1, x_2, x_3) must satisfy both of these inequalities, so it must also satisfy their sum:

$$(y_1 + 3y_2)x_1 + (4y_1 - y_2)x_2 + y_2x_3 \leq 2y_1 + 4y_2.$$

Now suppose that each y_i is larger than the i th coefficient of the objective function:

$$y_1 + 3y_2 \geq 4, \quad 4y_1 - y_2 \geq 1, \quad y_2 \geq 3.$$

This assumption lets us derive an upper bound on the objective value of *any* feasible solution:

$$4x_1 + x_2 + 3x_3 \leq (y_1 + 3y_2)x_1 + (4y_1 - y_2)x_2 + y_2x_3 \leq 2y_1 + 4y_2. \quad (*)$$

In particular, by plugging in the optimal solution (x_1^*, x_2^*, x_3^*) for the original LP, we obtain the following upper bound on σ^* :

$$\sigma^* = 4x_1^* + x_2^* + 3x_3^* \leq 2y_1 + 4y_2.$$

Now it's natural to ask how tight we can make this upper bound. How small can we make the expression $2y_1 + 4y_2$ without violating any of the inequalities we used to prove the upper bound? This is just another linear programming problem.

$$\begin{aligned} & \text{minimize} && 2y_1 + 4y_2 \\ & \text{subject to} && y_1 + 3y_2 \geq 4 \\ & && 4y_1 - y_2 \geq 1 \\ & && y_2 \geq 3 \\ & && y_1, y_2 \geq 0 \end{aligned}$$

In fact, this is precisely the dual of our original linear program! Moreover, inequality (*) is just an instantiation of the Weak Duality Theorem.

25.6 Strong Duality

The Fundamental Theorem can be rephrased in the following form:

Strong Duality Theorem. *If x^* is an optimal solution for a canonical linear program Π , then there is an optimal solution y^* for its dual Π^* , such that $c \cdot x^* = y^*Ax^* = y^* \cdot b$.*

Proof (sketch): I'll prove the theorem only for *non-degenerate* linear programs, in which (a) the optimal solution (if one exists) is a unique vertex of the feasible region, and (b) at most d constraint hyperplanes pass through any point. These non-degeneracy assumptions are relatively easy to enforce in practice and can be removed from the proof at the expense of some technical detail. I will also prove the theorem only for the case $n \geq d$; the argument for under-constrained LPs is similar (if not simpler).

To develop some intuition, let's first consider the *very* special case where $x^* = (0, 0, \dots, 0)$. Let e_i denote the i th standard basis vector, whose i th coordinate is 1 and all other coordinates are 0. Because $x_i^* = 0$ for all i , our non-degeneracy assumption implies the *strict* inequality $a_i \cdot x^* < b_i$ for all i . Thus, any sufficiently small ball around the origin does not intersect any other constraint hyperplane $a_i \cdot x = b_i$. Thus, for all i , and for any sufficiently small $\delta > 0$, the vector δe_i is feasible. Because x^* is the unique optimum, we must have $\delta c_i = c \cdot (\delta e_i) < c \cdot x^* = 0$. We conclude that $c_i < 0$ for all i .

Now let $y = (0, 0, \dots, 0)$ as well. We immediately observe that $yA \geq c$ and $y \geq 0$; in other words, y is a *feasible* solution for the dual linear program Π . But $y \cdot b = 0 = c \cdot x^*$, so the weak duality theorem implies that y is an *optimal* solution to Π , and the proof is complete for this very special case.

Now let us consider the more general case. Let x^* be the optimal solution for the linear program Π ; our non-degeneracy assumption implies that this solution is unique, and that exactly d of the n linear constraints are satisfied with equality. Without loss of generality (by permuting the constraints and possibly changing coordinates), we can assume that these are the first d constraints. Thus, we have

$$\begin{aligned} a_i \cdot x^* &= b_i && \text{for all } i \leq d, \\ a_i \cdot x^* &< b_i && \text{for all } i \geq d + 1, \end{aligned}$$

where a_i denotes the i th row of A . Let A_\bullet denote the $d \times d$ matrix containing the first d rows of A . Our non-degeneracy assumption implies that A_\bullet has full rank, and thus has a well-defined inverse $V = A_\bullet^{-1}$.

Now define a vector $y \in \mathbb{R}^n$ by setting

$$\begin{aligned} y_j &:= c \cdot v^j && \text{for all } j \leq d, \\ y_j &:= 0 && \text{for all } j \geq d + 1, \end{aligned}$$

where v^j denotes the j th column of $V = A_\bullet^{-1}$. Note that $a_i \cdot v^j = 0$ if $i \neq j$, and $a_i \cdot v^j = 1$ if $i = j$.

To simplify notation, let $y_\bullet = (y_1, y_2, \dots, y_d)$ and let $b_\bullet = (b_1, b_2, \dots, b_d) = A_\bullet x^*$. Because $y_i = 0$ for all $i \geq d + 1$, we immediately have

$$y \cdot b = y_\bullet \cdot b_\bullet = cVb_\bullet = cA_\bullet^{-1}b_\bullet = c \cdot x^*$$

and

$$yA = y_\bullet A_\bullet = cVA_\bullet = cA_\bullet^{-1}A_\bullet = c.$$

The point x^* lies on exactly d constraint hyperplanes; moreover, any sufficiently small ball around x^* intersects *only* those d constraint hyperplanes. Consider the point $\tilde{x} = x^* - \varepsilon v^j$, for some index $1 \leq j \leq d$ and some sufficiently small $\varepsilon > 0$. We have $a_i \cdot \tilde{x} = a_i \cdot x^* - \varepsilon(a_i \cdot v^j) = b_i$ for all $i \neq j$, and $a_j \cdot \tilde{x} = a_j \cdot x^* - \varepsilon(a_j \cdot v^j) = b_j - \varepsilon < b_j$. Thus, \tilde{x} is a feasible point for Π . Because x^* is the *unique* optimum for Π , we must have $c \cdot \tilde{x} = c \cdot x^* - \varepsilon(c \cdot v^j) < c \cdot x^*$. We conclude that $y_j = c \cdot v^j > 0$ for all j .

We have shown that $yA \geq c$ and $y \geq 0$, so y is a *feasible* solution for the dual linear program Π . We have also shown that $y \cdot b = c \cdot x^*$, so by the Weak Duality Theorem, y is also an *optimal* solution for Π , and the proof is complete! \square

We can also give a useful geometric interpretation to the vector $y_\bullet \in \mathbb{R}^d$. Each linear equation $a_i \cdot x = b_i$ defines a hyperplane in \mathbb{R}^d with normal vector a_i . The normal vectors a_1, \dots, a_d are linearly independent (by non-degeneracy) and therefore describe a coordinate frame for the vector space \mathbb{R}^d . The definition of y_\bullet implies that $c = y_\bullet A_\bullet = \sum_{i=1}^d y_i a_i$. In other words, **y_\bullet lists the coefficients of the objective vector c in the coordinate frame a_1, \dots, a_d .**

25.7 Complementary Slackness

Complementary Slackness Theorem. Let x^* be an optimal solution to a canonical linear program Π , and let y^* be an optimal solution to its dual Π . Then for every index i , we have $y_i^* > 0$ if and only if $a_i \cdot x^* = b_i$. Symmetrically, for every index j , we have $x_j^* > 0$ if and only if $y^* \cdot a^j = c_j$.

To be written

Exercises

- (a) Describe how to transform any linear program written in general form into an equivalent linear program written in slack form.

$$\begin{array}{l}
 \text{maximize } \sum_{j=1}^d c_j x_j \\
 \text{subject to } \sum_{j=1}^d a_{ij} x_j \leq b_i \quad \text{for each } i = 1 \dots p \\
 \sum_{j=1}^d a_{ij} x_j = b_i \quad \text{for each } i = p+1 \dots p+q \\
 \sum_{j=1}^d a_{ij} x_j \geq b_i \quad \text{for each } i = p+q+1 \dots n
 \end{array}
 \implies
 \begin{array}{l}
 \max \quad c \cdot x \\
 \text{s.t. } Ax = b \\
 x \geq 0
 \end{array}$$

- Describe precisely how to dualize a linear program written in slack form.
- Describe precisely how to dualize a linear program written in general form:

In all cases, keep the number of variables in the resulting linear program as small as possible.

- A matrix $A = (a_{ij})$ is *skew-symmetric* if and only if $a_{ji} = -a_{ij}$ for all indices $i \neq j$; in particular, every skew-symmetric matrix is square. A canonical linear program $\max\{c \cdot x \mid Ax \leq b; x \geq 0\}$ is *self-dual* if the matrix A is skew-symmetric and the objective vector c is *equal* to the constraint vector b .
 - Prove that any self-dual linear program Π is syntactically equivalent to its dual program Π .
 - Show that any linear program Π with d variables and n constraints can be transformed into a self-dual linear program with $n+d$ variables and $n+d$ constraints. The optimal solution to the self-dual program should include both the optimal solution for Π (in d of the variables) and the optimal solution for the dual program Π (in the other n variables).
- Give a linear-programming formulation of the *maximum-cardinality bipartite matching* problem. The input is a bipartite graph $G = (U \cup V; E)$, where $E \subseteq U \times V$; the output is the largest matching in G . Your linear program should have one variable for each edge.
 - Now dualize the linear program from part (a). What do the dual variables represent? What does the objective function represent? What problem is this!?
- Give a linear-programming formulation of the *minimum-cost feasible circulation problem*. Here you are given a flow network whose edges have both capacities and costs, and your goal is to find a feasible circulation (flow with value 0) whose cost is as small as possible.

5. An *integer program* is a linear program with the additional constraint that the variables must take only integer values.
- Prove that deciding whether an integer program has a feasible solution is NP-complete.
 - Prove that finding the optimal feasible solution to an integer program is NP-hard.

[Hint: Almost any NP-hard decision problem can be formulated as an integer program. Pick your favorite.]

- *6. *Helly's theorem* states that for any collection of convex bodies in \mathbb{R}^d , if every $d + 1$ of them intersect, then there is a point lying in the intersection of all of them. Prove Helly's theorem for the special case where the convex bodies are halfspaces. Equivalently, show that if a system of linear inequalities $Ax \leq b$ does not have a solution, then we can select $d + 1$ of the inequalities such that the resulting subsystem also does not have a solution. [Hint: Construct a dual LP from the system by choosing a 0 cost vector.]

7. Given points $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$ in the plane, the *linear regression problem* asks for real numbers a and b such that the line $y = ax + b$ fits the points as closely as possible, according to some criterion. The most common fit criterion is minimizing the L_2 error, defined as follows:⁷

$$\varepsilon_2(a, b) = \sum_{i=1}^n (y_i - ax_i - b)^2.$$

But there are several other fit criteria, some of which can be optimized via linear programming.

- (a) The L_1 error (or *total absolute deviation*) of the line $y = ax + b$ is defined as follows:

$$\varepsilon_1(a, b) = \sum_{i=1}^n |y_i - ax_i - b|.$$

Describe a linear program whose solution (a, b) describes the line with minimum L_1 error.

- (b) The L_∞ error (or *maximum absolute deviation*) of the line $y = ax + b$ is defined as follows:

$$\varepsilon_\infty(a, b) = \max_{i=1}^n |y_i - ax_i - b|.$$

Describe a linear program whose solution (a, b) describes the line with minimum L_∞ error.

⁷This measure is also known as *sum of squared residuals*, and the algorithm to compute the best fit is normally called (*ordinary/linear*) *least squares fitting*.