1 The linear algebra of linear programs (November 9, 2022)

Many optimization problems can be formulated as linear programs. The main features of a linear program are the following:

- Variables are *real numbers*. That is, they are continuous rather than discrete.

- The objective function is a linear function of the variables. (Each variable effects the objective function linearly, at a slope independent of the values of the other variables.)

- Constraints on the variables are linear.

As we shall see, there are polynomial time algorithms that find the optimal solution for linear programs. The versatility of linear programming makes it a very useful general purpose algorithmic tool. Linear programming, which can be viewed as a relaxation of integer programming, is also often used in order to find approximate solutions to NP-hard combinatorial optimization problems. We shall touch upon this subject as well. The theory of linear programming provides insights that allow us to prove theorems unrelated to algorithmic questions – an aspect that will also be demonstrated in the course. The course will also consider extensions of linear programming, such as semidefinite programming, and specializations, such as algorithms that apply to specific types of linear programs.

The course will be accompanied by lecture notes. In addition there are many books on linear programming, including [2] which is freely accessible on the web.

1.1 The diet example

One of the first recorded examples of a linear program is the so called *diet* problem. Assume that a person's diet is composed entirely of n foods. Assume that there are m nutrients, and the person is required to consume at least b_i units of nutrient i (for $1 \leq i \leq m$). Let a_{ij} denote the amount of nutrient i present in one unit of food j. Let c_i denote the cost of one unit of food item i. The problem is to design a diet of minimal cost that supplies at least the required amount of nutrients.

Using matrix and vector representation, we have the following linear program:

minimize $c^T x$
subject to
$Ax \geq b$
$x \geq 0$

Let us remark that one could add constraints limiting from above the amount of nutrients consumed (e.g., not too much fat in the diet).

(Modeling the diet problem as a linear program includes many hidden assumptions such as that the cost is linear, that there are no interdependencies in the nutritional values of various foods, that any fractional amount of food can be consumed, etc. As these assumptions are only approximations of the true state of affairs, an optimal solution to the linear program can only be viewed as a guideline for selecting a desirable solution, and not as the ultimate solution.)

A solution satisfying all constraints is *feasible*. A feasible solution that also optimizes the objective function is *optimal*.

1.2 Fitting a polynomial to noisy data

Suppose that we believe that there is a degree d polynomial p that determines the relation between two variables x and y. Namely, $y = p(x) = \sum_{j=0}^{d} a_j x^j$. We wish to find the coefficients of this polynomial. For this we sample pairs (x_i, y_i) of data points. For example, the x variable may denote inverse of distance from transmitter, and the y variable may denote the signal intensity, and we perform measurements of signal intensities at various distances from the transmitter. If measurements are exact, then recovering the unknown coefficients of the polynomial (namely, *interpolation*) can be done by solving a system of linear equations. Every data point (x_i, y_i) gives one equation $y_i = \sum_{j=0}^{d} a_j (x_i)^j$ (where y_i and x_i are given, and the a_j are the variables). Given $d+1$ data points (with distinct x_i), the $d+1$ constraints can be arranged as the rows of a so called *Vandermonde* matrix V, giving the system of linear equations $Va = y$, where a is the column vector $(a_0, \ldots, a_d)$ of unknown coefficients that we wish to find, and y is the known column vector $(y_0, \ldots, y_d)$. As Vandermonde matrices are invertible, we obtain $a = V^{-1} y$.

Suppose now a more realistic scenario in which there is some error in measuring the y_i values. In this case, there might not be any degree d polynomial that fits the data. Once there are more than $d+1$ data points, the system of linear equations might not be consistent. Now the task is to search for the degree d polynomial that fits the data best. One way of measuring how well a polynomial fits the data is via a min-max criterion – we wish to minimize ϵ so that for every point x, the value of the polynomial is within ϵ of the corresponding value of y.

We can express the above problem as a linear program whose variables are ϵ and $a_0, \ldots, a_d$. The number of data points is $m \geq d+1$. Arrange the constraints implied by the data points in rows of an m by $d+1$ matrix A, where $A_{ij} = (x_i)^j$. The vectors a and y are as above, and 1^m denotes the all 1 vector.

minimize ϵ
subject to
$Aa \geq y - \epsilon 1^m$
$Aa \leq y + \epsilon 1^m$
$\epsilon \geq 0$

1.3 Fourier-Motzkin elimination

In the coming lectures we will see approaches for solving linear programs. But first, let us sketch an inefficient approach.

Consider first the case of finding a feasible solution. If all constraints were equalities rather than inequalities, we could use Gaussian elimination (and find a feasible solution in polynomial time). For inequalities we may perform a similar (but less efficient) procedure known as Fourier-Motzkin elimination. To eliminate a variable x_i, we bring all inequalities in which it participates into one of two forms: $x_i \leq \alpha_j$ and $x_i \geq \beta_j$, where for every $1 \leq j \leq m$, α_j and β_j are linear functions of the rest of the variables. We can now eliminate x_i, and replace all constraints that included it by the new constraints $\beta_j \leq \alpha_k$ (for $1 \leq j \leq m$, $1 \leq k \leq m$). Repeating this procedure for the other variables gives eventually a set of inequalities in one variable, that can easily be solved. Note that each variable elimination at most squares the number of constraints. Hence the complexity of the above procedure

is $O(m^{2^n})$ in the worst case. (Note however that if we can in each elimination step choose a variable that appears in at most 4 inequalities, then the number of inequalities does not grow, and the procedure becomes polynomial.)

Given the ability to find feasible solutions, we can in principle also find near optimal solutions. This can be done by replacing the objective function by a constraint $c^T x \leq t$, and searching for feasible solution. By doing binary search on t (assuming that we have reasonable estimates for lower and upper bounds of t), we can find a solution arbitrarily close to the true optimum. For the specific case of using Fourier-Motzkin elimination, we can treat t as a variable and leave it last in the elimination order. This allows us to find the exact optimum.

We shall now develop some theory that will allow us to design more efficient algorithms for solving linear programs.

1.4 Forms

The linear program for the diet problem is in *canonical form*. In a linear program in *general form*, the constraints are linear but may involve inequalities of both types ($\leq$ and $\geq$), as well as equalities ($=$). Variables may be required to be nonnegative ≥ 0, or else be unconstrained. Another useful form of a linear program is the *standard form*:

minimize $c^T x$

subject to

$Ax = b$

$x \geq 0$

All forms are equivalent in terms of their expressive power, and it is simple to transform a linear program in general form to standard form and to canonical form.

To replace the objective to maximization, change it to **maximize** $-c^T x$.

To replace an unconstrained variable x_i by nonnegative variables, introduce two auxiliary variables x_i^+ and x_i^-, replace every occurrence of x_i by $x_i^+ - x_i^-$, and add the two nonnegativity constraints $x_i^+ \geq 0$ and $x_i^- \geq 0$.

To change a $\leq$ inequality $A_i^T x \leq b_i$ to a $\geq$ inequality, negate it, obtaining $-A_i^T x \geq -b_i$.

To change an equality $A_i^T x = b_i$ to inequalities, replace it by $A_i^T x \leq b_i$ and $A_i^T x \geq b_i$.

To change an inequality $A_i^T x \leq b_i$ to an equality, introduce a *slack variable* y_i, the inequality $A_i^T x + y_i = b_i$, and the nonnegativity constraint $y_i \geq 0$.

Note that the above transformations change a general linear program with n variables and m constraints to a linear program with $O(n + m)$ variables and constraints.

For linear programs in standard form, it is convenient to assume that the constraints (rows of the matrix A) are linearly independent. If the rows are not linearly independent, then it suffices to consider rows of A that constitute a basis for the row space (a maximal linearly independent set of row vectors). Either every solution that satisfies the constraints that correspond to the basis satisfies all constraints, or the LP is infeasible.

1.5 LP versus systems of linear equations

The feasibility question of a linear program modifies the question of solving a system of linear equations in two different aspects. For linear programs in canonical form, if we incorporate the nonnegativity constraints into the main set of constraints, we see that

the difference compared to systems of linear equations is that we have a system of linear inequalities, rather than equalities. For linear programs in standard form, we have a system of linear equations $Ax = b$, and the difference compared to systems of linear equations is that we seek only nonnegative solutions (satisfying $x \geq 0$) to this system. Though the above outlines two different modifications to systems of linear equations, both of them offer the same kind of algorithmic challenges, because of the equivalence between different forms of linear programs.

1.6 Basic feasible solutions

Consider an LP in standard form, with linearly independent constraints. Let B be a sub-matrix of A containing exactly m linearly independent columns. This is a *basis* of the column space of A. Let x_B be the set of basic variables corresponding to the columns of B. If $B^{-1}b \geq 0$, then the following is a basic feasible solution: the basic variables are set to $B^{-1}b$, and the nonbasic variables are set to 0. Clearly this solution is feasible. Note that it satisfies n linearly independent constraints with equality: the m constrains of $Ax = b$, and $n - m$ of the nonnegativity constraints. The other (nonnegativity) constraints are also satisfied, though not necessarily with equality.

Each basis gives at most one basic feasible solution. (It gives none if the condition $B^{-1}b \geq 0$ fails to hold.) Two different bases may give the same basic feasible solution, in which case the basic feasible solution is *degenerate* (more than $n - m$ variables are set to 0).

Basic feasible solutions play an important role in linear programs.

Lemma 1 *For every basic feasible solution, there is a choice of objective function $c^t x$ for which the bfs is the unique optimal solution.*

Proof: Consider an objective function for which $c_i = 0$ if x_i is positive in the bfs, and $c_i = 1$ otherwise. The objective function is nonnegative for every feasible solution (because of the nonnegativity constraints), and 0 for the particular bfs. As there is no other feasible solution with the same (or larger by containment) set of 0 variables (otherwise $B^{-1}b$ would have more than one solution), the bfs is the unique minimum. $\square$

Lemma 2 *Every LP in standard form is either infeasible, or the optimal value is unbounded, or it has a basic feasible solution that is optimal.*

Proof: Assume that the LP is feasible and that its optimum is bounded from below. Use q to denote the value of the objective function. First we show that for every value of q, if there is a solution x of value q (namely, $c^t x = q$), then there is a basic feasible solution x^* of value at most q.

Let x^+ be the set of positive variables in x. If the columns in A corresponding to x^+ are linearly independent, then we can complete them to a basis giving x as a bfs. Hence it remains to deal with the case that the columns in A corresponding to x^+ are not linearly independent. Denote them by $A^1, \ldots, A^t$ and the values of the positive variables by $z_1, \ldots, z_t$. Then we have $\sum_{i=1}^{t} A^i z_i = b$. There is also a nonzero solution w to $\sum_{i=1}^{t} A^i x_i = 0$. Consider now the sign of $c^t w$. If $c^t w < 0$, then perturbing x by a positive multiple of w reduces the value of the objective function. As the value of the objective function is bounded from below, it must be the case that changing x by a multiple of w eventually causes a new

nonnegativity constraint to become tight. This gives a feasible solution with lower value for the objective function and at least one more variable that is 0. If $c^t w > 0$, then the same effect is reached by perturbing x by a negative multiple of w. If $c^t w = 0$ then using the fact that w has at least one nonzero variable (and is 0 whenever x is), there is a multiple of w that we can add to x that causes one more variable to become 0, without changing the value of the objective function. Repeating the above process we must eventually reach some solution x^* for which the columns in A corresponding to $(x^*)^+$ are linearly independent, implying a bfs.

As there are only finitely many basic feasible solutions (at most $\binom{n}{m}$), there is one with minimum value, and this must be an optimum solution for the LP. $\quad \square$

Recall Cramer's rule for solving $Bx = b$, where B is an invertible order n matrix. The solution is

$$x_j = \frac{\det B^{j \leftarrow b}}{\det B}$$

for $1 \leq j \leq n$, where here $B^{j \leftarrow b}$ is the matrix B with column j replaced by b. If each entry in B and b is an integer with absolute value at most M, then each x_j is a rational number with numerator and denominator bounded by at most $M^n n!$. This can be used to show that the length of numbers involved in a basic feasible solution is polynomially related to the input size. (Moreover, it can be shown that when a system of linear equations is solved by Gaussian elimination, the length of intermediate numbers produced by the algorithm is also polynomially related to the input size.)

Lemma 2 and its proof have the following algorithmic consequences for LPs in standard form:

1. Given a feasible solution of value q, one can find in polynomial time a basic feasible solution of value at most q.

2. In order to solve an LP optimally, it suffices to consider only basic feasible solutions. As there are at most $\binom{n}{m}$ basic feasible solutions, we can solve LPs optimally in this time.

One can extend the notion of basic feasible solutions also to LPs that are not in standard form, provided that among the constraints of the LP there are n constraints that are linearly independent. (This applies for example to LPs in canonical form, because the n nonnegativity constraints are linearly independent.) In this case, a basic feasible solution is one that satisfies n linearly independent constraints with equality. Moreover, for every LP that has n constraints that are linearly independent, if it is feasible and bounded, then it has a bfs that is optimal. (Proving this statement is left as a homework assignment.)

1.7 An application – the Beck-Fiala theorem

The existence of basic feasible solutions is a property that is used in many contexts. Here is an instructive example.

There is a set S on n items. We are given a collection of subsets $S_i \subset S$, $1 \leq i \leq m$. The goal is to color the items red and blue, such that every subset is "nearly balanced", namely, has roughly the same number of red items as blue items. The (absolute value of

the) maximum difference between number of red and blue items in a subset is called the *discrepancy* of the coloring. The Beck-Fiala theorem [1] gives sufficient conditions for a coloring of low discrepancy to exist, regardless of the values of n and m.

Theorem 3 *If every item is contained in an most d subsets, then there is a coloring with discrepancy at most $2d - 1$.*

Proof: Think of 0 as denoting the color red and 1 as denoting the color blue. Think of x_j as a variable representing the color of item j. Let a_{ij} be 1 if set i contains item j, and 0 otherwise. Hence we seek a 0/1 solution satisfying $|\sum_j a_{ij}x_j - |S_i|/2| \leq d - 1/2$ for every set S_i.

Consider the following set of linear constraints:

1. For all $1 \leq j \leq n$, $x_j \geq 0$.

2. For all $1 \leq j \leq n$, $x_j \leq 1$.

3. Discrepancy constraints: For all $1 \leq i \leq m$, $\sum_{1 \leq j \leq n} a_{ij}x_j = |S_i|/2$.

This set has a feasible solution, namely, $x_j = 1/2$. We shall "round" this fractional solution to an integer solution. This rounding will introduce a slackness of at most $d - 1/2$ in the discrepancy constraints.

As the set of constraints is feasible, it has a basic feasible solution. Let $m_I \leq m$ be the number of linearly independent discrepancy constraints. In a bfs n linearly independent constraints are satisfied with equality, and hence at least $n - m_I$ variables are set to either 0 or 1. Update the set of constraints by removing the integer variables, and updating the right hand side of every discrepancy constraint by the sum of values removed from its left hand side. Some discrepancy constraints are "small", in the sense that they have at most d noninteger variables left. For them, no matter how we set the remaining variables, the slackness introduced will be at most $d - 1/2$. (We used here the fact that the right hand side of each constraint is half-integer, and the right hand side of small constraints must lie between $1/2$ and $d - 1/2$.) Hence we can ignore the small constraints. So we are left with a set of residual constraints and a set of variables whose fractional values satisfy the residual constraints. As each residual discrepancy constraint has more than d variables, and each variable appears in at most d residual constraints, then m', the number of residual discrepancy constraints, must be smaller than n', the number of fractional variables. Again, moving to a bfs of the residual system of constraints, $n' - m'_I > 0$ variables receive integral value. Continuing with this process until no residual constraints are left, the theorem is proved. $\square$

We note that the proof (together with the proof of Lemma 2) gives a polynomial time algorithm for finding a coloring of low discrepancy.

1.8 An open question

It is quite easy to slightly strengthen the Beck-Fiala theorem. In the proof, call a set small if it has at most $d - 1$ fractional variables. Then for small sets the slackness will be at most $d - 3/2$, giving a discrepancy of at most $2d - 3$. Now it might not be the case that

all sets eventually end up being small – all residual sets might contain exactly d fractional variables. In this case, rounding all remaining fractional variables to their nearest integer gives a slackness of at most $d/2$. Hence the discrepancy is at most $\max[2d - 3, d]$. For d equals 1 and 2, this is best possible. (For $d = 1$, a single set with one item has discrepancy 1. For $d = 2$, three sets, each containing two out of three items, requires discrepancy 2.) For $d \geq 3$ we have that $2d - 3 \geq d$, and hence the discrepancy is at most $2d - 3$, a mild improvement over the Beck-Fiala theorem.

It is conjectured that under the terms of the Beck-Fiala theorem, there always is a coloring with discrepancy $O(\sqrt{d})$, offering a large improvement when d is large. If true, this bound would be best possible (up to constant multiplicative factors), as even for a collection of d random subsets of a set of d items, (with high probability) every coloring has a discrepancy of $\Omega(\sqrt{d})$. (This last statement can be proven by a straightforward probabilistic argument.)

References

[1] Jozsef Beck and Tibor Fiala (1981). "Integer-making" theorems. Discrete Applied Mathematics. 3 (1): 1-8.

[2] Jiri Matousek, Bernd Gartner: Understanding and Using Linear Programming. Springer 2007. http://link.springer.com/book/10.1007%2F978-3-540-30717-4