CS168: The Modern Algorithmic Toolbox
Lecture #18: Linear and Convex Programming, with Applications to Sparse Recovery

Tim Roughgarden & Gregory Valiant*

May 31, 2024

1 Linear Programming

1.1 Context

The more general a problem, the more computationally difficult it is. For example, sufficient generalization of a polynomial-time solvable problem often yields an NP-hard problem. If you only remember one thing about linear programming, make it this: *linear programming is a remarkable sweet spot balancing generality and computational tractability*, arguably more so than any other problem in the entire computational landscape.

Zillions of problems, including ℓ_1-minimization, reduce to linear programming. It would take an entire course to cover even just its most famous applications. Some of these applications are conceptually a bit boring but still very important — as early as the 1940s, the military was using linear programming to figure out the most efficient way to ship supplies from factories to where they were needed. Central problems in computer science that reduce to linear programming include maximum flow and bipartite matching. (There are also specialized algorithms for these two problems, see CS261.) Linear programming is also useful for NP-hard problems, for which it serves as a powerful subroutine in the design of heuristics (again, see CS261).

Despite this generality, linear programs can be solved efficiently, both in theory (meaning in worst-case polynomial time) and in practice (with input sizes up into the millions).

1.2 Using Linear Programming

You can think of linear programming as a restricted programming language for encoding computational problems. The language is flexible, and sometimes figuring out the right way to use it requires some ingenuity (as we'll see).

At a high level, the description of a linear program specifies what's allowed, and what you want. Here are the ingredients:

1. *Decision variables.* These are real-valued variables $x_1, \ldots, x_n \in \mathbb{R}$. They are "free," in the sense that it is the job of the linear programming solver to figure out the best joint values for these variables.

2. *Constraints.* Each constraint should be linear, meaning it should have the form

$$\sum_{j=1}^{n} a_{ij} x_j \leq b_i$$

   or

$$\sum_{j=1}^{n} a_{ij} x_j = b_i.$$

   We didn't bother including constraints of the form $\sum_{j=1}^{n} a_{ij} x_j \geq b_i$, since these are equivalent to $\sum_{j=1}^{n} (-a_{ij}) x_j \leq -b_i$. All of the a_{ij}'s and b_i's are real-valued constants, meaning specific numbers (1, -5, 10, etc.) that are hard-coded into the linear program.

3. *Objective function.* Again, this should be linear, of the form

$$\min \sum_{j=1}^{n} c_j x_j.$$

   It's fine to maximize instead of minimize: after all, $\max \sum_{j=1}^{n} c_j x_j$ yields the same result as $\min \sum_{j=1}^{n} (-c_j) x_j$.

So what's not allowed in a linear program? Terms like x_j^2, $x_j x_k$, $\log(1 + x_j)$, etc. So whenever a decision variable appears in an expression, it is alone, possibly multiplied by a constant. These linearity requirements may seem restrictive, but many real-world problems are well approximated by linear programs.

1.3 A Simple Example

To make linear programs more concrete and develop your intuition about them, let's look at a simple example. Suppose there are two decision variables x_1 and x_2 — so we can visualize solutions as points (x_1, x_2) in the plane. See Figure 1. Let's consider the (linear) objective function of maximizing the sum of the decision variables:

$$\max\ x_1 + x_2. \tag{1}$$

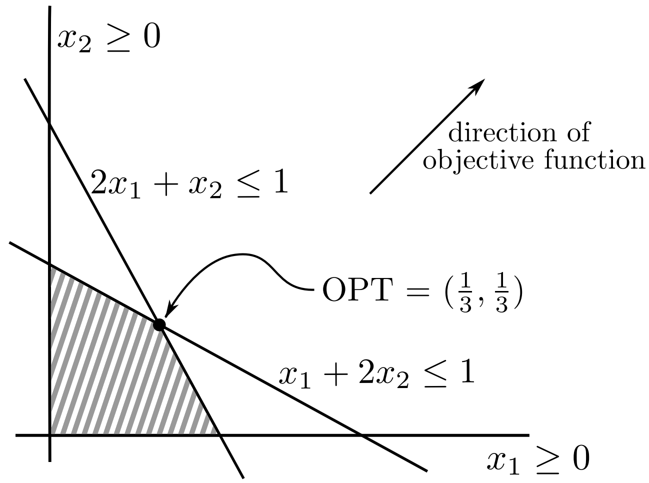

Figure 1: A linear program in 2 dimensions.

We'll look at four (linear) constraints:

$$x_1 \geq 0 \quad (2)$$
$$x_2 \geq 0 \quad (3)$$
$$2x_1 + x_2 \leq 1 \quad (4)$$
$$x_1 + 2x_2 \leq 1. \quad (5)$$

The first two inequalities restrict feasible solutions to the non-negative quadrant of the plane. The second two inequalities further restrict feasible solutions to lie in the shaded region depicted in Figure 1. Geometrically, the objective function asks for the feasible point furthest in the direction of the coefficient vector $(1, 1)$ — the "most northeastern" feasible point. Eyeballing the feasible region, this point is $(\frac{1}{3}, \frac{1}{3})$, for an optimal objective function value of $\frac{2}{3}$.

1.4 Geometric Intuition

This geometric picture remains valid for general linear programs, with an arbitrary number of dimensions and constraints: *the objective function gives the optimization direction, and the goal is to find the feasible point that is furthest in this direction.* Moreover, the feasible region of a linear program is just a higher-dimensional analog of a polygon.[1]

1.5 Algorithms for Linear Programming

Linear programs are not difficult to solve in two dimensions — for example, one can just check all of the vertices (i.e., "corners") of the feasible region. In high dimensions, linear

[1] Called a "polyhedron;" in the common special case where the feasible region is bounded, it is called a "polytope."

programs are not so easy; the number of vertices can grow exponentially with the number of dimensions (e.g., think about hypercubes), so there's no time to check them all. Nevertheless, we have the following important fact.

Fact 1.1 *Linear programs can be solved efficiently.*

The theoretical version of Fact 1.1 states that there is a polynomial-time algorithm for linear programming.[2] The practical version of Fact 1.1 is that there are excellent commercial codes available for solving linear programs.[3] These codes routinely solve linear programs with millions of variables and constraints. One thing to remember about linear programming is that, for over 60 years, many people with significant resources — ranging from the military to large companies — have had strong incentives to develop good codes for it. This is one of the reasons that the best codes are so fast and robust.

There are a variety of ways to efficiently solve linear programs. In lecture we briefly discussed the high level idea for using the multiplicative weights algorithm (learning with expert advice) to solve linear programs. I'll try to add the details of that to these lecture notes at some point soon.

2 Linear Programming and ℓ_1-Minimization

We now show that the ℓ_1-minimization problem in compressive sensing can be solved using linear programming. The only non-trivial issue is the objective function

$$\min \|\mathbf{x}\|_1 = \sum_{j=1}^{n} |x_j|,$$

which, because of the absolute values, is non-linear.

As a warm-up, suppose first that we know that the unknown signal $\mathbf{z}$ is component-wise non-negative (in addition to being k-sparse). Then, the ℓ_1-minimization problem is just

$$\min \sum_{j=1} x_j$$

subject to

$$\mathbf{A}\mathbf{x} = \mathbf{b} \tag{6}$$

and

$$\mathbf{x} \geq 0. \tag{7}$$

The objective function is clearly linear. The n non-negativity constraints in (7) — each of the form $x_j \geq 0$ for some j — are linear. Each of the m equality constraints (6) has the form $\sum_{j=1}^{n} a_{ij} x_j = b_i$, and is therefore linear. Thus, this is a linear program.

[2] The earliest, from 1979, is the "ellipsoid method" [7]; this was a big enough deal at the time that it made the New York Times [1].

[3] The open-source solvers are not as good, unfortunately, but are still useful for solving reasonably large linear programs.

We know the unknown signal $\mathbf{z}$ satisfies $\mathbf{Ax} = \mathbf{b}$ (by the definition of $\mathbf{b}$). We're also assuming that $\mathbf{z} \geq 0$. Hence, $\mathbf{z}$ is a feasible solution to the linear program. Since $\mathbf{x} \geq 0$ for every feasible solution, the objective function value $\sum_{j=1}^{n} x_j$ equals $\|\mathbf{x}\|_1$ for every feasible solution. We conclude that this linear program is a faithful encoding of ℓ_1-minimization for non-negative signals.

For the general case of real-valued signals $\mathbf{z}$, the key trick is to add additional variables that allow us to "linearize" the non-linear objective function. In addition to the previous decision variables $x_1, \ldots, x_n$, our linear program will include auxiliary decision variables $y_1, \ldots, y_n$. The intent is for y_j to represent $|x_j|$. We use the objective function

$$\min \sum_{j=1}^{n} y_j, \tag{8}$$

which is clearly linear. We also add $2n$ linear inequalities, of the form

$$y_j - x_j \geq 0 \tag{9}$$

and

$$y_j + x_j \geq 0 \tag{10}$$

for every $j = 1, 2, \ldots, n$. Finally, we have the usual m linear consistency constraints

$$\mathbf{Ax} = \mathbf{b}. \tag{11}$$

Every feasible solution of this linear program satisfies all of the constraints, and in particular (9) and (10) imply that $y_j \geq \max\{x_j, -x_j\} = |x_j|$ for every $j = 1, 2, \ldots, n$. Observe further that at an optimal solution, equality must hold for every j: given a feasible solution with $y_j > x_j$ and $y_j > -x_j$ for some j, one can decrease y_j slightly to produce a new solution that is still feasible and that has slightly better (i.e., smaller) objective function value (8). It follows that the values of the variables $\mathbf{x}$ in an optimal solution to the linear program given by (8)–(11) is the optimal solution to the ℓ_1-minimization problem.

To further showcase the power and flexibility of linear programming, suppose that the results of the linear measurements are corrupted by noise. Concretely, assume that instead of receiving $b_i = \langle \mathbf{a}_i, \mathbf{z} \rangle$ for each measurement $i = 1, 2, \ldots, m$, we receive a value $b_i \in [\langle \mathbf{a}_i, \mathbf{z} \rangle - \epsilon, \langle \mathbf{a}_i, \mathbf{z} \rangle + \epsilon]$, where $\epsilon > 0$ is a bound on the magnitude of the noise. Now, the linear system $\mathbf{Ax} = \mathbf{b}$ might well be infeasible — $\mathbf{z}$ is now only an approximately feasible solution. The linear program (8)–(11) is easily modified to accommodate noise — just replace the equality constraints (11) by two sets of inequality constraints,

$$\sum_{j=1}^{n} a_{ij} x_j \leq b_i + \epsilon$$

and

$$\sum_{j=1}^{n} a_{ij} x_j \geq b_i - \epsilon$$

Figure 2: Examples of convex and non-convex sets.

for each $i = 1, 2, \ldots, m$. The guarantee in of recovering sparse signals via ℓ_1 minimization can also be extended, with significant work, to handle noise [4].

Remark 2.1 (Keep Linear Programming in Your Toolbox) This concludes our brief discussion of linear programming. While compressive sensing is a convenient excuse to discuss this powerful tool, don't forget that linear programming is useful for solving or approximating a huge range of applications drawn from many different domains. It's quite likely that one or more problems arising in your future work will be solvable using linear programming. The classic book [5] remains an excellent introduction to some of the applications.

3 Beyond Linear Programs: Convexity

We next discuss a generalization of linear programming that captures still more applications, without sacrificing too much computational efficiency.

3.1 Convex Sets

A good rule of thumb is to equate "convex" with "nice" and "non-convex" with "nasty," especially when optimization is concerned. This rule of thumb holds for simple algorithms, like gradient descent: when minimizing convex functions, gradient descent has nice properties, including that gradient descent will find the global minimizer, whereas for non-convex functions gradient descent might only find a local optima. Here, convexity is in large part what's driving the computational tractability of linear programming.

Convexity is relevant for both sets and for functions. Intuitively, a subset $C \subseteq \mathbb{R}^n$ is convex if it is "filled in," meaning that it contains all line segments between its points. See Figure 2 for examples. Formally, C is *convex* if for every $\mathbf{x}, \mathbf{y} \in C$ and $\lambda \in [0, 1]$, $\lambda\mathbf{x} + (1 - \lambda)\mathbf{y} \in C$. (As λ ranges from 0 to 1, it traces out the line segment from $\mathbf{y}$ to $\mathbf{x}$.)

For example, the feasible region of every linear program is convex. To see this, first suppose there is only one constraint, which is an inequality. Then the feasible region is just a half-space, which is clearly convex. The feasible region of a linear program is an intersection of such half-spaces. (Note that an equality constraint is equivalent to the combination of two inequality constraints.) The intersection of convex sets C_1, C_2 is again convex — if $\mathbf{x}$

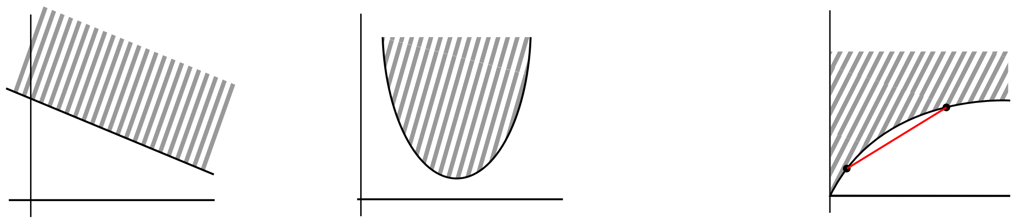

$f(x) = ax + b$ $f(x) = $ quadratic $f(x) = \sqrt{x}$

convex non-convex

Figure 3: Examples of convex and non-convex functions.

and $\mathbf{y}$ are in both C_1 and C_2, then the line segment between $\mathbf{x}$ and $\mathbf{y}$ lies inside both C_1 and C_2 (since each is convex), so this line segment also lies in their intersection. We conclude that every linear program has a convex feasible region.

For a relevant example that is more general than the finite intersection of half-spaces and subspaces, take C to be the set of $n \times n$ symmetric and positive semidefinite (PSD) matrices, viewed as a subset of $\mathbb{R}^{n^2}$.[4] It is clear that the set of symmetric matrices is convex — the average of symmetric matrices is again symmetric. It is true but less obvious that the set remains convex under the extra PSD constraint.[5]

3.2 Convex Functions

Who had the nerve to use the same word "convex" for two totally different things, sets and functions? The overloaded terminology becomes more forgivable if we define a function $f : \mathbb{R}^n \to \mathbb{R}$ to be *convex* if and only if the region above its graph is a convex set. See Figure 3 for some examples.

Equivalently, a convex function is one where all "chords" of its graph lie above the graph. Mathematically, this translates to

$$f(\lambda\mathbf{x} + (1-\lambda)\mathbf{y}) \le \lambda f(\mathbf{x}) + (1-\lambda) f(\mathbf{y})$$

for every $\mathbf{x}, \mathbf{y} \in C$ and $\lambda \in [0,1]$. That is, for points $\mathbf{x}$ and $\mathbf{y}$, if you take the average of $\mathbf{x}$ and $\mathbf{y}$ and then apply f, you'll get a smaller number than if you first apply f to $\mathbf{x}$ and $\mathbf{y}$ and then average the results. It's not always easy to check whether or not a given function is convex, but there is a mature analytical toolbox for this purpose (taught in EE364, for example).

[4] There are many equivalent definitions of PSD matrices. One of the simplest is as the matrices of the form $\mathbf{A}^T\mathbf{A}$, like the covariance matrices we were looking at during our PCA discussions in Lectures #7–9.

[5] Another definition of PSD matrices is as the matrices $\mathbf{A}$ for which the corresponding quadratic form $\mathbf{x}^T\mathbf{A}\mathbf{x}$ is nonnegative for every $\mathbf{x} \in \mathbb{R}^n$. Using linearity, it is easy to see that the average of two matrices that satisfy this condition yields another matrix that satisfies the condition.

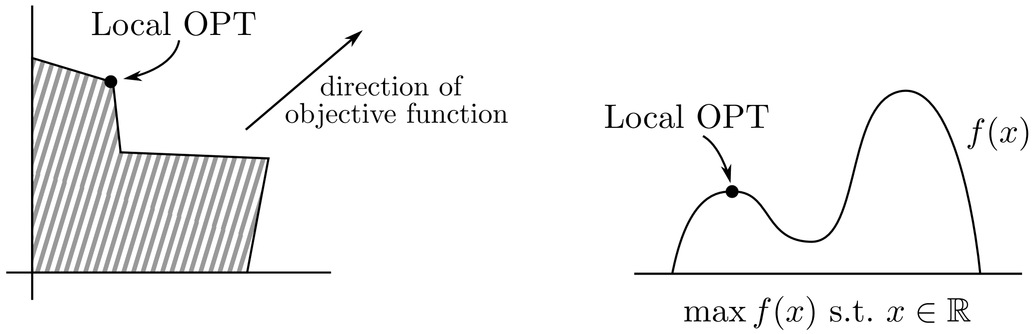

Figure 4: Non-convexity and local optima. (Left) A linear (i.e. convex) objective function with a non-convex feasible region. (Right) A non-convex objective function over a convex feasible region (the real line).

3.3 Convex Programs

Convexity leads to computational tractability. For example, in theory, it is possible to minimize an essentially arbitrary convex function over an essentially arbitrary convex feasible region. (There's a bit of fine print, but the conditions are quite mild.) This is fantastic news: in principle, we should be able to develop fast and robust algorithms for all of the convex optimization problems that we want to solve.

Practice is in the process of catching up with what the theory predicts. To oversimplify the current state-of-the-art, there are currently solvers that can handle medium-size and sufficiently nice convex optimization problems. The first piece of good news is that this is already enough to solve many problems that we're interested in. The second piece of good news is that, as we speak, many smart people are working hard to close the gap in computational efficiency between linear and convex programming solvers.

Summarizing: convex programming is even more general than linear programming and captures some extra interesting applications. It is relatively computationally tractable, although the biggest instance sizes that can be solved are generally one or two orders of magnitude smaller than with linear programming (e.g., tens of thousands instead of millions).

Remark 3.1 (Why Convexity Helps) For intuition about why convexity leads to tractability, consider the case where the feasible region or the objective function is *not* convex. With a non-convex feasible region, there can be "locally optimal" feasible points that are not globally optimal, even with a linear objective function (Figure 4(left)). The same problem arises with a non-convex objective function, even when the feasible region is just the real line (Figure 4(right)). When both the objective function and feasible region are convex, this can't happen — all local optima are also global optima. As you might expect, this makes optimization much easier.

References

[1] M. W. Browne. A Soviet discovery rocks world of mathematics. *New York Times*, November 7 1979.

[2] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[3] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.

[4] E. J. Candes, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.

[5] G. W. Dantzig. *Linear Programming and Extensions.* Princeton, 1963.

[6] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.

[7] L. G. Khachiyan. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady*, 20(1):191–194, 1979.

[8] J. Matousek and B. Gärtner. *Understanding and Using Linear Programming.* Springer, 2006.

[9] B. Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12:3413–3430, 2011.