

## Examples of constrained convex optimization problems

We will close our initial discussion of constrained convex optimization with a very brief tour of common categories of constrained convex optimization problems, giving a few practical examples where each arises. This discussion is by no means exhaustive, but is merely intended to help you to have some concrete examples in the back of your mind where the techniques we have developed can be applied.

### Linear programming

Perhaps the simplest constrained convex optimization problem to write down (although not necessarily the easiest to solve) is a **linear program** (LP). An LP minimizes a linear objective function subject to multiple linear constraints:

$$\underset{\mathbf{x}}{\text{minimize}} \quad \mathbf{c}^T \mathbf{x} \quad \text{subject to} \quad \mathbf{a}_m^T \mathbf{x} \leq b_m, \quad m = 1, \dots, M.$$

The general form above can include linear equality constraints  $\mathbf{a}_i^T \mathbf{x} = b_i$  by enforcing both  $\mathbf{a}_i^T \mathbf{x} \leq b_i$  and  $(-\mathbf{a}_i)^T \mathbf{x} \leq b_i$ . We can also write the  $M$  constraints compactly as  $\mathbf{Ax} \leq \mathbf{b}$ , where  $\mathbf{A}$  is the  $M \times N$  matrix with the  $\mathbf{a}_m^T$  as rows.

Linear programs do not necessarily have to have a solution; it is possible that there is no  $\mathbf{x}$  such that  $\mathbf{Ax} \leq \mathbf{b}$ , or that the program is unbounded in that there exists a series  $\mathbf{x}_1, \mathbf{x}_2, \dots$ , all obeying  $\mathbf{Ax}_k \leq \mathbf{b}$ , with  $\lim \mathbf{c}^T \mathbf{x}_k \rightarrow -\infty$ .

There is no formula for the solution of a general linear program. Fortunately, there exists very reliable and efficient software for solving

them. The first LP solver was developed in the late 1940s (Dantzig’s “simplex algorithm”, which is a clever iterative descent algorithm tailored to the LP setting), and now LP solvers are considered a mature technology. If the constraint matrix  $\mathbf{A}$  is structured, then linear programs with millions of variables can be solved to high accuracy on a standard computer.

Linear programs are a very important class of optimization problems. However, if a single constraint (or the objective function) are nonlinear, then we move into the much broader class of **nonlinear programs**, which has really been the main focus of this course.

## Example: Chebyshev approximations

Suppose that we want to find the vector  $\mathbf{x}$  so that  $\mathbf{Ax}$  does not vary too much in its maximum deviation:

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad \max_{m=1, \dots, M} |y_m - \mathbf{a}_m^T \mathbf{x}| = \underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad \|\mathbf{y} - \mathbf{Ax}\|_\infty.$$

This is called the **Chebyshev approximation problem**.

We can solve this problem with linear programming. To do this, we introduce the auxiliary variable  $u \in \mathbb{R}$  — it should be easy to see that the program above is equivalent to

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}^N, u \in \mathbb{R}}{\text{minimize}} \quad u \quad \text{subject to} \quad & y_m - \mathbf{a}_m^T \mathbf{x} \leq u \\ & y_m - \mathbf{a}_m^T \mathbf{x} \geq -u \\ & m = 1, \dots, M. \end{aligned}$$

To put this in the standard linear programming form, take

$$\mathbf{z} = \begin{bmatrix} \mathbf{x} \\ u \end{bmatrix}, \quad \mathbf{c}' = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}, \quad \mathbf{A}' = \begin{bmatrix} -\mathbf{A} & -1 \\ \mathbf{A} & -1 \end{bmatrix}, \quad \mathbf{b}' = \begin{bmatrix} -\mathbf{y} \\ \mathbf{y} \end{bmatrix},$$

and then solve

$$\underset{\mathbf{z} \in \mathbb{R}^{N+1}}{\text{minimize}} \quad \mathbf{c}'^T \mathbf{z} \quad \text{subject to} \quad \mathbf{A}' \mathbf{z} \leq \mathbf{b}'.$$

One natural application of this arises in the context of **filter design**. The standard “filter synthesis” problem is to find an finite-impulse response (FIR) filter whose discrete-time Fourier transform (DTFT) is as close as possible to some desired  $H^*(\omega)$ . If we measure “closeness” in terms of the maximum deviation, the result is called an “equiripple design” since the error in the solution will tend to have ripples a uniform distance away from the ideal.

We will not go through the derivation in detail here, but I will note that Georgia Tech’s very own Jim McClellan worked all of this out in the early 1970s with his advisor Tom Parks, developing the now ubiquitous Parks-McClellan filter design algorithm.

## Quadratic programming

Let us briefly return to our standard least squares problem. As we have seen before, this is equivalent to the problem of

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad \mathbf{x}^T \mathbf{A}^T \mathbf{A} \mathbf{x} - 2\mathbf{y}^T \mathbf{A} \mathbf{x}.$$

Suppose you now wanted to enforce some additional structure on  $\mathbf{x}$ . For example, you might have reason to desire a solution with only non-negative values. In adding such a constraint, we arrive at an example of a **quadratic program** (QP).

A QP minimizes a quadratic functional subject to linear constraints:

$$\underset{\mathbf{x}}{\text{minimize}} \quad \mathbf{x}^T \mathbf{H} \mathbf{x} + \mathbf{c}^T \mathbf{x}, \quad \text{subject to} \quad \mathbf{A} \mathbf{x} \leq \mathbf{b}.$$

If  $\mathbf{H}$  is symmetric positive semidefinite (i.e., symmetric with nonnegative eigenvalues), then the program is convex. If  $\mathbf{H}$  has even a single negative eigenvalue, then solving the program above is NP-hard.

QPs are almost as ubiquitous as LPs; they have been used in finance since the 1950s (see the example below), and are found all over operations research, control systems, and machine learning. As with LPs, there are reliable solvers and can be considered a mature technology.

A **quadratically constrained quadratic program** (QCQP) allows (convex) quadratic inequality constraints:

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & \mathbf{x}^T \mathbf{H} \mathbf{x} + \mathbf{c}^T \mathbf{x}, \quad \text{subject to} \quad \mathbf{x}^T \mathbf{H}_m \mathbf{x} + \mathbf{c}_m^T \mathbf{x} \leq b_m, \\ & m = 1, \dots, M. \end{aligned}$$

This program is convex if all of the  $\mathbf{H}_m$  are symmetric positive semidefinite; we are minimizing a convex quadratic functional over a region defined by an intersection of ellipsoids.

## Example: Portfolio optimization

One of the classic examples in convex optimization is finding investment strategies that “optimally”<sup>1</sup> balance the risk versus the return. The following quadratic program formulation is due to Markowitz, who formulated it in the 1950s, then won a Nobel Prize for it in 1990.

We want to spread our money over  $N$  different assets; the fraction of our money we invest in asset  $n$  is denoted  $x_n$ . We have the immediate

---

<sup>1</sup>I put “optimally” in quotes because, like everything in finance and the world, this technique finds the optimal answer for a specified model. The big question is then how good your model is ...

constraints that

$$\sum_{n=1}^N x_n = 1, \quad \text{and} \quad 0 \leq x_n \leq 1, \quad \text{for } n = 1, \dots, N.$$

The expected return on these investments, which are usually calculated using some kind of historical average, is  $\mu_1, \dots, \mu_N$ . The  $\mu_n$  are specified as multipliers, so  $\mu_n = 1.16$  means that asset  $n$  has a historical return of 16%. We specify some target expected return  $\rho$ , which means

$$\sum_{n=1}^N \mu_n x_n \geq \rho.$$

We want to solve for the  $\mathbf{x}$  that achieves this level of return while minimizing our *risk*. Here, the definition of risk is simply the variance of our return — if the assets have covariance matrix  $\mathbf{R}$ , then the risk of a given portfolio allocation  $\mathbf{x}$  is

$$\text{Risk}(\mathbf{x}) = \mathbf{x}^T \mathbf{R} \mathbf{x} = \sum_{m=1}^M \sum_{n=1}^M R_{mn} x_m x_n.$$

Our optimization program is then<sup>2</sup>

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \mathbf{x}^T \mathbf{R} \mathbf{x} \\ & \text{subject to} && \boldsymbol{\mu}^T \mathbf{x} \geq \rho \\ & && \mathbf{1}^T \mathbf{x} = 1 \\ & && \mathbf{0} \leq \mathbf{x} \leq \mathbf{1}. \end{aligned}$$

This is an example of a QP with linear constraints. It is convex since the matrix  $\mathbf{R}$  is a covariance matrix, and so by construction it is symmetric positive semidefinite.

---

<sup>2</sup>Throughout these notes, we will use  $\mathbf{1}$  for a vector of all ones, and  $\mathbf{0}$  for a vector of all zeros.

# Example: Support vector machines

As we have already seen, SVMs are a classical approach for designing a classifier in machine learning and can be expressed as

$$\underset{\mathbf{w}, b}{\text{minimize}} \quad \|\mathbf{w}\|_2^2 \quad \text{subject to} \quad (\mathbf{x}_i^T \mathbf{w} + b)y_i \geq 1 \text{ for all } i.$$


This is another example of a QP with linear constraints.

## Second-order cone programs

A **second-order cone program** (SOCP) is an optimization problem where the constraint set forms what is called, perhaps unsurprisingly, a second-order cone. The canonical example of a second-order cone is the set:

$$\{(\mathbf{x}, t), \mathbf{x} \in \mathbb{R}^N, t \in \mathbb{R} : \|\mathbf{x}\|_2 \leq t\}.$$

This is a subset of  $\mathbb{R}^{N+1}$ . Here is an example in  $\mathbb{R}^3$ :



The standard form of a SOCP is

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \mathbf{c}^T \mathbf{x} \\ & \text{subject to} && \|\mathbf{A}_m \mathbf{x} + \mathbf{b}_m\|_2 \leq \mathbf{c}_m^T \mathbf{x} + d_m, \quad m = 1, \dots, M. \end{aligned}$$

We have a linear objective function and constraints that require  $(\mathbf{y}, t)$  to lie inside the second-order cone, where  $\mathbf{y}$  and  $t$  are allowed to be any affine function of  $\mathbf{x}$ .

SOCPs turn out to be much more common than you might initially expect. First, it is not hard to show that an LP is also a SOCP. It turns out that QPs and (convex) QCQPs are also SOCPs, so we can think of SOCPs as a generalization of what we have already seen. However, the class of possible SOCPs also includes many optimization problems beyond what we have seen so far.

## Example: Generalized geometric medians

Suppose that we have  $M$  points  $\mathbf{p}_1, \dots, \mathbf{p}_M \in \mathbb{R}^N$  and that we would like to find the “center” of this set of points. The *geometric median* is the point  $\mathbf{x}$  that minimizes the sum (or equivalently, average) of the distances to the points  $\mathbf{p}_1, \dots, \mathbf{p}_M$ . This can be posed as the optimization problem

$$\underset{\mathbf{x}}{\text{minimize}} \quad \sum_{m=1}^M \|\mathbf{x} - \mathbf{p}_m\|_2.$$

In the case where  $N = 1$ , this is equivalent to the standard median. The special case of  $M = 3$  points in a dimension  $N \geq 2$  was first considered by Pierre de Fermat, with Evangelista Torricelli providing a simple geometric solution in the 17<sup>th</sup> century. In general, however, there is no closed-form solution to this problem.

It is relatively straightforward to show that this problem can be cast as a SOCP. Specifically, it should be clear that it is equivalent to:

$$\begin{aligned} & \underset{\mathbf{x}, t}{\text{minimize}} && \sum_{m=1}^M t_m \\ & \text{subject to} && \|\mathbf{x} - \mathbf{p}_m\|_2 \leq t_m, \quad m = 1, \dots, M. \end{aligned}$$

A slight variation on this problem is to try to minimize the maximum distance from  $\mathbf{x}$  to the  $\mathbf{p}_m$ :

$$\underset{\mathbf{x}}{\text{minimize}} \quad \max_{m \in \{1, \dots, M\}} \|\mathbf{x} - \mathbf{p}_m\|_2.$$

This too has a simple formulation as a SOCP:

$$\begin{aligned} & \underset{\mathbf{x}, t}{\text{minimize}} && t \\ & \text{subject to} && \|\mathbf{x} - \mathbf{p}_m\|_2 \leq t, \quad m = 1, \dots, M. \end{aligned}$$

## Semidefinite programs

So far we have typically been looking at problems where we are optimizing over vectors  $\mathbf{x} \in \mathbb{R}^N$ . In many important applications, our decision variables are more naturally represented as a matrix  $\mathbf{X}$ . In such problems, it is common to encounter the constraint that this matrix  $\mathbf{X}$  must be positive semidefinite. When the objective function is linear and we have affine constraints, this is called a **semidefinite program** (SDP).

To state the standard form for an SDP, it is useful to introduce some notation. First, we will let  $\mathcal{S}^N$  denote the set of  $N \times N$  symmetric



matrices, and  $\mathcal{S}_+^N$  the set of symmetric positive semidefinite matrices. Furthermore, we let

$$\langle \mathbf{Y}, \mathbf{X} \rangle = \text{trace}(\mathbf{Y}^T \mathbf{X})$$

denote the (trace) inner product between a pair of matrices.<sup>3</sup> With this notation in hand, the standard form for an SDP is given by

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} && \langle \mathbf{C}, \mathbf{X} \rangle \\ & \text{subject to} && \langle \mathbf{A}_m, \mathbf{X} \rangle \leq b_m, \quad m = 1, \dots, M \\ & && \mathbf{X} \in \mathcal{S}_+^N, \end{aligned}$$

where  $\mathbf{C}, \mathbf{A}_1, \dots, \mathbf{A}_M \in \mathcal{S}$ .

SDPs are the broadest class of convex problems that we will consider in this course. All of the problems we have looked at so far (LPs, QPs, SOCPs) can be shown to be special cases of SDPs.

## Example: Bounding portfolio risk

Let us briefly return to our previous example of portfolio optimization. Before we assumed that we knew the expected returns and the covariance matrix  $\mathbf{R}$  for the different assets under consideration, and our goal was to determine the optimal allocation. Here we consider a slightly different problem. Suppose that we already have a fixed allocation  $\mathbf{x}$  across the different assets, but rather than knowing the covariance matrix  $\mathbf{R}$  exactly, we assume that we have only an estimate of  $\mathbf{R}$ . A natural question is whether we can quantify how large the true risk of our portfolio might be in such a case.

---

<sup>3</sup>This is simply the inner product that would result from reshaping  $\mathbf{X}$  and  $\mathbf{Y}$  into vectors and applying the standard inner product.

Suppose that we have confidence intervals on how accurate our covariance estimate is of the form

$$L_{mn} \leq R_{mn} \leq U_{mn}.$$

For a given portfolio  $\mathbf{x}$ , we can compute the maximum possible risk of that portfolio that is consistent with the given bounds via the following SDP:

$$\begin{aligned} & \underset{\mathbf{R}}{\text{maximize}} && \mathbf{x}^T \mathbf{R} \mathbf{x} \\ & \text{subject to} && L_{mn} \leq R_{mn} \leq U_{mn}, \quad m, n = 1, \dots, N \\ & && \mathbf{R} \in \mathcal{S}_+^N. \end{aligned}$$

We have to enforce the constraint that  $\mathbf{R} \in \mathcal{S}_+^N$  because  $\mathbf{R}$  must be a covariance matrix, and ignoring this constraint would yield a risk that is not actually achievable.