

Lecture 9

Conic optimization

Instructor: Prof. Gabriele Farina

In Lectures 3, we considered normal cones at the intersection of linear equality and inequality constraints. In Exercise 1 of Homework 1, we were also able to give closed formulas for the normal cone in certain sets, such as a ball. In Lecture 7, we turned our attention to more general feasible sets. There, we saw how KKT conditions define necessary conditions for feasible sets defined by *functional constraints* by approximating the feasible set with the linearization of the active constraints and considering the normal cone to the intersection of those linearizations.

Today, we look at a class of convex feasible sets that are neither linear equality nor inequality constraints. Also, they are not defined via functional constraints, making the KKT machinery inapplicable.

L9.1 Conic optimization

A *conic* optimization problem is a nonlinear optimization problem whose feasible set is the intersection between an affine subspace (that is, a system of linear equalities) and a nonempty closed convex cone \mathcal{K} :

$$\begin{aligned} \min_x & f(x) \\ \text{s.t.} & Ax = b \\ & x \in \mathcal{K} \end{aligned}$$

This class of problems captures linear programming, semidefinite programming, second-order cone programming, copositive programming, and more, depending on the specific cone \mathcal{K} that is selected.

L9.1.1 Nonnegative cone \leftrightarrow Linear programming

The first example is the nonnegative cone $\mathbb{R}_{\geq 0}^n$. In this case, the conic problem takes the more familiar form

$$\begin{aligned} \min_x & f(x) \\ \text{s.t.} & Ax = b \\ & x \geq 0 \end{aligned}$$

*These notes are class material that has not undergone formal peer review. The TAs and I are grateful for any reports of typos.

which is a *linear program*. So, for the specific choice of the nonnegative cone, conic programming corresponds to linear programming.

L9.1.2 Ice-cream (aka. Lorentz) cone ↔ Second-order conic programming


Definition L9.1.

The *ice-cream cone*, or *Lorentz cone*, is defined as

$$\mathcal{L}^n := \{(x, z) \in \mathbb{R}^{n-1} \times \mathbb{R} : z \geq \|x\|_2\}.$$

The figure on the right shows the shape of \mathcal{L}^3 .

Conic problems for the specific choice of the Lorentz cone are usually called *second-order cone programs*, or *SOCP* for short. Several problems of interest in engineering and physics can be modeled as SOCP, including the design of antenna arrays, the positions of spring systems at rest, and grasp planning in robotics.



L9.1.3 Semidefinite cone ↔ Semidefinite programming


Definition L9.2. The *semidefinite cone* \mathcal{S}^n is the set of positive semidefinite $n \times n$ matrices:

$$\begin{aligned} \mathcal{S}^n &:= \{X \in \mathbb{S}^n : X \succeq 0\} \\ &= \{X \in \mathbb{S}^n : a^\top X a \geq 0 \quad \forall a \in \mathbb{R}^n\}, \end{aligned}$$

where \mathbb{S}^n is the set of symmetric $n \times n$ real matrices.

The figure on the right shows the cone of values (x, y, z) corresponding to the semidefinite cone of 2×2 matrices

$$\begin{pmatrix} x & y \\ y & z \end{pmatrix} \succeq 0.$$



In this simple 2×2 case, the positive semidefiniteness condition is equivalent to

$$\begin{aligned} x + z &\geq 0 && \text{(the sum of the eigenvalues [trace] is nonnegative)} \\ xz - y^2 &\geq 0 && \text{(the product of the eigenvalues [determinant] is nonnegative.)} \end{aligned}$$

These conditions are equivalent to $x, z \geq 0 \wedge xz \geq y^2$. So, slices of the cone with planes orthogonal to the z -axis are shaped like parabolas $x \geq y^2/z$. Similarly, slices with planes orthogonal to the x -axis are shaped like parabolas $z \geq y^2/x$. Note how the curvature of the parabolas decreases as the slicing plane gets further away from the origin.

Remark L9.1. Semidefinite programming subsumes both linear programming and second-order cone programming. This is because

$$x \in \mathbb{R}_{\geq 0}^n \iff \begin{pmatrix} x_1 & & \\ & \ddots & \\ & & x_n \end{pmatrix} \succeq 0, \text{ and}$$

$$z \geq \|x\|_2 \iff \begin{pmatrix} zI & x \\ x^\top & z \end{pmatrix} \succeq 0.$$

Semidefinite programming is an extremely powerful tool with applications in all disciplines. I have noted resources that treat semidefinite programming extensively at the end of this document; we will see connections to optimization of polynomials in the next lecture.

L9.1.4 Copositive cone \leftrightarrow Copositive programming


Definition L9.3. The *copositive cone* \mathcal{C}^n is the set of symmetric $n \times n$ real matrices:

$$\mathcal{C}^n := \{X \in \mathbb{S}^n : a^\top X a \geq 0 \quad \forall a \in \mathbb{R}_{\geq 0}^n\}.$$

The difference with the positive semidefinite cone \mathcal{S}^n lies in the fact that we need $a^\top X a \geq 0$ *only for nonnegative vectors* $a \in \mathbb{R}_{\geq 0}^n$, instead of all $a \in \mathbb{R}^n$.

The figure on the right shows the cone of values (x, y, z) corresponding to the copositive cone of 2×2 matrices, that is

$$\begin{pmatrix} x & y \\ y & z \end{pmatrix} \in \mathcal{C}^2.$$



It follows immediately from the definition that \mathcal{C}^n is a superset of both the cone of nonnegative symmetric matrices $\mathcal{N}^n := \mathbb{S}^n \cap \mathbb{R}_{\geq 0}^{n \times n}$ and the cone of semidefinite matrices. Hence,

$$\mathcal{N}^n + \mathcal{S}^n \subseteq \mathcal{C}^n \quad \text{for all } n \in \mathbb{N}.$$

Curiously, the reverse inequality is known to hold, *but only up to dimension* $n \leq 4$ [Dia62]:

$$\mathcal{C}^n = \mathcal{S}^n + \mathcal{N}^n \quad \text{if } n \leq 4.$$

Beyond dimension 4, some matrices are copositive, and yet they are not expressible as the sum of \mathcal{S}^n with \mathcal{N}^n .

Remark L9.2. One important fact to know about the copositive cone is that despite being a closed convex cone, optimization of even linear functions over the copositive cone is computationally intractable!

In fact, even checking if a matrix M belongs to \mathcal{C}^n is in general NP-hard.

For example, the size of the *maximum independent set* of a simple graph $G = (V, E)$, that is, the size of the largest set $S \subseteq V$ of vertices with the property that no two vertices in S are connected by an edge in E , can be computed as the solution to the following copositive problem. Let $A \in \{0, 1\}^{V \times V}$ denote the adjacency matrix of G , that is,

$$A_{ij} = \begin{cases} 1 & \text{if } (i, j) \in E \\ 0 & \text{otherwise.} \end{cases}$$

The size of the maximum independent set of G can then be found by solving the copositive optimization problem

$$\begin{aligned} \min_{w, t \in \mathbb{R}, Q} \quad & w \\ \text{s.t.} \quad & Q = wI + tA - \mathbf{1} \mathbf{1}^\top \\ & Q \in \mathcal{C}^n \end{aligned} \tag{MIS}$$

where $\mathbf{1} \in \mathbb{R}^V$ is the vector of all ones. More information is available in the paper by [de Klerk, E., & Pasechnik, D. V. \[dP02\]](#). Since deciding whether a graph admits an independent set of size k cannot take polynomial time unless $\mathbf{P} = \mathbf{NP}$, this means that solving copositive optimization problems must take more than polynomial time in the worst case unless $\mathbf{P} = \mathbf{NP}$. The hardness applies already for linear objectives, as the previous example shows. So, we conclude—for example—that one cannot construct an efficient separation oracle for \mathcal{C}^n .

Several techniques are known to “relax” the copositive cone by, for example, approximating it with a (higher-dimensional) semidefinite cone. The monographs of [Gärtner, B., & Matoušek, J. \[GM12\]](#) and [Klerk, E. de. \[Kle02\]](#) do an excellent job at explaining approximation techniques that involve semidefinite cones.

L9.2 Optimality conditions and conic duality

In general, the feasible set Ω of conic optimization problems cannot be easily captured via *functional* constraints. Because of that, we cannot use the KKT conditions we discussed in the previous lecture to compute the normal cone at each point of the feasible set. In this situation, we need to compute the normal cone from first principles.

To simplify the treatment and avoid the machinery of *relative* interiors (interiors in the topology induced by the smallest affine subspace that contains a set), we will assume in the following that \mathcal{K} is a cone with interior. This condition is satisfied by all the cases discussed above.

In order to compute the normal cone at a point at the intersection of $\{Ax = b\}$ and \mathcal{K} , we will use the following general result (again based on separation), which enables us to consider the affine subspace and the cone separately and sum their normal cones.

Theorem L9.1. Let $H := \{x \in \mathbb{R}^n : Ax = b, \text{ with } A \in \mathbb{R}^{m \times n}\}$ be an affine subspace, and S be a closed convex set (not necessarily a cone) such that $S^\circ \cap H$ is nonempty, where S° is the interior of S . For any $x \in H \cap S$,

$$\mathcal{N}_{H \cap S}(x) = \mathcal{N}_H(x) + \mathcal{N}_S(x).$$

So, using the fact (see Lecture 2) that $\mathcal{N}_H(x) = \text{colspan}(A^\top)$, we have

$$\mathcal{N}_{H \cap S}(x) = \{A^\top \mu + z : \mu \in \mathbb{R}^m, z \in \mathcal{N}_S(x)\}.$$

Remark L9.3. The condition of existence of a feasible point in the interior is sometimes referred to as *Slater's condition*, *strict feasibility condition*, or *strong feasibility condition*. The term *Slater's condition* is appropriate in light of the constraint qualification condition for the KKT theory we discussed in Lecture 7. In both cases, the insistence on an interior solution is required to rule out pathological behavior.

We know from Lectures 3 that Slater's condition was not required for the intersection of halfspaces, which are flat surfaces. The condition becomes, however, crucial when considering sets that can have curvature. For example, consider the intersection of a two-dimensional ball of radius 1 centered in $(x, y) = (1, 0)$ with the plane $x = 0$. The normal cone at the intersection is \mathbb{R}^2 , but this is not equal to the sum between the normal cones of the sets at the intersection points (which evaluates to $\mathbb{R} \times \{0\}$).

Remark L9.4. The previous result can be generalized, without much effort nor substantially new ideas to the case of intersection between two generic convex sets whose relative interiors have a nonempty intersection. In other words, the *normal cone to the intersection between convex sets is equal to the sum of the normal cones to the individual sets* if the relative interiors of the two sets intersect. This also explains why the strict feasibility condition imposed by Slater is sufficient.

L9.2.1 The normal cone at a point in \mathcal{K}

With Theorem L9.1 in our toolbox, the only obstacle to writing first-order optimality conditions for conic problems is being able to compute the normal at points in \mathcal{K} . Luckily, we can do that in closed form in all the abovementioned cases.

In particular, it turns out that the normal cone at any point in \mathcal{K} can be defined starting from the normal cone at 0 (remember that 0 is always part of a nonempty cone, as a straightforward consequence of the definition of cone). The normal cone for a cone \mathcal{K} at $0 \in \mathcal{K}$ is such an important object, that several names for it or its related quantities are used in the literature:

- The *polar cone* of \mathcal{K} , often denoted \mathcal{K}^\perp , is exactly the normal cone at 0:

$$\mathcal{K}^\perp := \mathcal{N}_{\mathcal{K}}(0) = \{d : \langle d, y \rangle \leq 0 \quad \forall y \in \mathcal{K}\}.$$

- The *dual cone* of \mathcal{K} , often denoted \mathcal{K}^* , is the opposite of the normal cone at 0:

$$\mathcal{K}^* := -\mathcal{K}^\perp = -\mathcal{N}_{\mathcal{K}}(0) = \{d : \langle d, y \rangle \geq 0 \quad \forall y \in \mathcal{K}\}.$$

For all the important cones considered earlier, we have a very precise idea of what their polar and dual cones look like:

Theorem L9.2.

1. The nonnegative cone, the Lorentz cone, and the semidefinite cones are *self-dual cones*, that is, $\mathcal{K}^* = \mathcal{K}$.
2. The dual cone to the copositive cone \mathcal{C}^n is called the *totally positive cone*, defined as

$$\mathcal{P}^n := \left\{ \sum_{i=1}^k z_i z_i^\top : z_i \in \mathbb{R}_{\geq 0}^n, k \in \mathbb{N} \right\} = \{BB^\top : B \in \mathbb{R}_{\geq 0}^{n \times m}, m \in \mathbb{N}\}$$

[▷ Try to prove the previous theorem.] Since for a closed convex cone \mathcal{K} , $(\mathcal{K}^*)^* = \mathcal{K}$, it follows that the dual to the totally positive cone is the copositive cone. [▷ You should also try to prove that $(\mathcal{K}^*)^* = \mathcal{K}$.]

Once the normal cone at 0 is established, the normal cone at any point $x \in \mathcal{K}$ can be recovered as a function of $\mathcal{N}_{\mathcal{K}}(0)$.

Theorem L9.3. Let $\mathcal{K} \subseteq \mathbb{R}^n$ be a nonempty, closed convex cone. The normal cone at any point $x \in \mathcal{K}$ is given by

$$\mathcal{N}_{\mathcal{K}}(x) = \{d \in \mathcal{K}^\perp : \langle d, x \rangle = 0\}.$$

Proof. Fix a generic $x \in \mathcal{K}$. We break the proof into a proof of the two separate inclusions.

(\subseteq) Let $d \in \mathbb{R}^n$ be a direction in $\mathcal{N}_{\mathcal{K}}(x)$, that is,

$$\langle d, y - x \rangle \leq 0 \quad \forall y \in \mathcal{K}.$$

Since \mathcal{K} is a cone, we have that both $2x$ and $\frac{1}{2}x$ belong to \mathcal{K} . So,

$$\begin{aligned} \langle d, 2x - x \rangle \leq 0 &\implies \langle d, x \rangle \leq 0, \text{ and} \\ \left\langle d, \frac{1}{2}x - x \right\rangle \leq 0 &\implies \langle d, x \rangle \geq 0, \end{aligned}$$

showing that $\langle d, x \rangle = 0$ is necessary. But then, from the condition that $\langle d, y - x \rangle \leq 0$ for all $y \in \mathcal{K}$ we deduce that $\langle d, y \rangle \leq 0$ for all $y \in \mathcal{K}$, which means that $d \in \mathcal{K}^\perp = \mathcal{N}_{\mathcal{K}}(0)$.

(\supseteq) Vice versa, let $d \in \{d \in \mathcal{K}^\perp : \langle d, x \rangle = 0\}$. Then, for any $y \in \mathcal{K}$,

$$\begin{aligned} \langle d, y - x \rangle &= \langle d, y \rangle - \langle d, x \rangle \\ &= \langle d, y \rangle && \text{(since } \langle d, x \rangle = 0\text{)} \\ &\leq 0 && \text{(since } d \in \mathcal{K}^\perp\text{)}. \end{aligned}$$

□

L9.2.2 First-order optimality conditions

We now have all the ingredients to write first-order optimality conditions for a conic problem

$$\begin{aligned} \min_x & f(x) \\ \text{s.t.} & Ax = b \\ & x \in \mathcal{K}. \end{aligned} \tag{2}$$

Specifically, by combining Theorem L9.1 with Theorem L9.3 and the sufficiency of first-order optimality conditions seen in Lecture 3, we obtain the following result.

Theorem L9.4. If f is convex and differentiable, and the conic optimization (2) is *strictly feasible*, that is, there exists $x \in \mathcal{K}^\circ$ such that $Ax = b$, then a point $x^* \in \Omega := \mathcal{K} \cap \{x : Ax = b\}$ is a minimizer of f on Ω if and only if

$$-\nabla f(x^*) \in \{A^\top \mu + z : \mu \in \mathbb{R}^m, z \in \mathcal{N}_{\mathcal{K}}(x^*)\},$$

that is, expanding $\mathcal{N}_{\mathcal{K}}(x^*)$ using Theorem L9.3, if and only if

$$\begin{aligned} -\nabla f(x^*) &\in \{A^\top \mu + z : \mu \in \mathbb{R}^m, z \in \mathcal{K}^\perp, \langle z, x^* \rangle = 0\} \\ \Leftrightarrow -\nabla f(x^*) &\in \{A^\top \mu - z : \mu \in \mathbb{R}^m, z \in \mathcal{K}^*, \langle z, x^* \rangle = 0\}. \end{aligned}$$

Remark L9.5. The condition $\langle z, x^* \rangle = 0$ is called the *conic complementary slackness condition*. It is a generalization of the complementary slackness condition we saw in Lecture 3 for linear programming. In the special case where $\mathcal{K} = \mathbb{R}_{\geq 0}^n$, the conic complementary slackness condition reduces to the complementary slackness condition for linear programming.

L9.2.3 Conic duality for linear objectives

In particular, when $f(x) = \langle c, x \rangle$ is a *linear* objective, then optimality of x^* for (2) is equivalent to the existence of μ, z such that

$$-c = A^\top \mu - z, \quad \text{where } \mu \in \mathbb{R}^m, z \in \mathcal{K}^*, \langle z, x^* \rangle = 0.$$

This suggests that whenever the conditions of strict feasibility holds, then a duality theory for conic problems can be derived following the exact same steps as we did in Lecture 3. In particular, we find that the dual problem is

$$\begin{aligned} \max_{\mu, z} & \langle b, \mu \rangle \\ \text{s.t.} & z = c - A^\top \mu \\ & \mu \in \mathbb{R}^m \\ & z \in \mathcal{K}^*. \end{aligned}$$

After showing that the value of the dual is always upper bounded by the value of the primal (weak duality) using the same steps as in Lecture 3, we conclude the following. [▷ Try working out the details, it should only take a minute.]

Theorem L9.5. If the primal problem

$$\begin{aligned} \min_x & \langle c, x \rangle \\ \text{s.t.} & Ax = b \\ & x \in \mathcal{K} \end{aligned}$$

is strictly feasible and admits an optimal solution x^* , then the dual problem

$$\begin{aligned} \max_{\mu, z} \quad & \langle b, \mu \rangle \\ \text{s.t.} \quad & z = c - A^\top \mu \\ & \mu \in \mathbb{R}^m \\ & z \in \mathcal{K}^* \end{aligned}$$

admits an optimal solution (μ^*, z^*) such that:

- the values of the two problems coincide: $\langle c, x^* \rangle = \langle b, \mu^* \rangle$; and
- the solution z^* satisfies the complementary slackness condition $\langle x^*, z^* \rangle = 0$.

► **Lagrangian duality in the conic case.** Alternatively, it is possible to adapt the Lagrangian ideas we saw in Lecture 8 (which, after all, were nothing but a rewrite of the first-order optimality conditions) to conic problems.

The key difference is that the constraint $x \in \mathcal{K}$ is not a functional constraint, so it cannot be easily included in the Lagrangian. For this, we consider some intermediate Lagrangian problem in which we only penalize violation of the constraint $Ax = b$, but otherwise retain the constraint $x \in \mathcal{K}$ as part of the domain. In other words, we consider the Lagrangian

$$\mathcal{L}(x; \mu) : \mathcal{K} \times \mathbb{R}^m, \quad \mathcal{L}(x; \mu) = \langle c, x \rangle - \langle \mu, Ax - b \rangle.$$

The conditions in Theorem L9.4 guarantee that for any solution x^* of (2), there exist $\mu^* \in \mathbb{R}^m$ such that

$$-c - A^\top \mu^* \in \{z \in \mathcal{K}^\perp : \langle z, x^* \rangle = 0\} = \mathcal{N}_{\mathcal{K}}(x^*).$$

and therefore x^* is a minimizer of $\mathcal{L}(x; \mu^*)$ over \mathcal{K} . Hence, using the same steps as in Lecture 8, we obtain the strong Lagrangian duality statement

$$\max_{\mu \in \mathbb{R}^m} \inf_{x \in \mathcal{K}} \mathcal{L}(x; \mu) = \langle c, x^* \rangle = \min_{x \in \mathcal{K}} \sup_{\mu \in \mathbb{R}^m} \mathcal{L}(x; \mu).$$

One can easily check that the Lagrangian dual problem on the left is the same as the dual problem we derived earlier in Theorem L9.5. Indeed,

$$\begin{aligned} \inf_{x \in \mathcal{K}} \mathcal{L}(x; \mu) &= \inf_{x \in \mathcal{K}} \langle c, x \rangle - \langle \mu, Ax - b \rangle \\ &= \langle \mu, b \rangle + \inf_{x \in \mathcal{K}} \langle x, c - A^\top \mu \rangle. \end{aligned}$$

The infimum on the right-hand side is equal to 0 if $(c - A^\top \mu) \in \mathcal{K}^*$, and $-\infty$ otherwise [► Show this!]. Hence, the dual problem is indeed equivalent to

$$\begin{aligned} \max_{\mu, z} \quad & \langle b, \mu \rangle \\ \text{s.t.} \quad & c - A^\top \mu \in \mathcal{K}^*. \end{aligned}$$

— L9.2.4 Failure of duality when the constraint qualification is not satisfied

It is essential to remember that duality might fail when the strict feasibility constraint qualification is not satisfied. Specifically, the primal problem might be feasible and yet the dual might not have an optimal solution.

The failure mode can be of different kinds:

- the primal has an optimal solution, the supremum of the dual matches the value of the primal, and yet the dual does not have a maximizer (Example L9.1); or
- the primal has an optimal solution, the dual has an optimal solution, but the values of the problems differ (Example L9.2).

Example L9.1 ([Kle02]). Consider the problem

$$\begin{aligned} \min_X \quad & 2X_{12} + X_{22} \\ \text{s.t.} \quad & X_{11} = 0 \\ & X_{22} = 1 \\ & X \succeq 0. \end{aligned}$$

The only point in the feasible set is $\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$; therefore, the optimal value is 1.

Since the matrix $\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$ is *not* in the interior of the semidefinite cone (one eigenvalue is 0), strict feasibility does not hold.

The dual problem in this case is

$$\begin{aligned} \max_{\lambda, Z} \quad & \lambda_2 \\ \text{s.t.} \quad & Z = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} - \lambda_1 \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} - \lambda_2 \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \\ & \lambda \in \mathbb{R}^2 \\ & Z \succeq 0 \end{aligned}$$

For $\lambda = (-\frac{1}{\epsilon}, 1 - \epsilon)$, $\epsilon > 0$ we obtain the feasible point $Z = \begin{pmatrix} 1/\epsilon & 1 \\ 1 & \epsilon \end{pmatrix} \succeq 0$ with objective value $1 - \epsilon$. Hence, the supremum of the dual matches the value of the primal, but the dual has no optimal solution.

Example L9.2 ([BV04]). Consider the primal problem

$$\begin{aligned} \min_X \quad & x_2 \\ \text{s.t.} \quad & X = \begin{pmatrix} x_2 + 1 & 0 & 0 \\ 0 & x_1 & x_2 \\ 0 & x_2 & 0 \end{pmatrix} \\ & X \succeq 0. \end{aligned}$$

Since every principal minor must be nonnegative, it follows immediately that $x_2 = 0$ for all feasible points, and so the objective value is 0. The dual problem is (after some manipulations)

$$\begin{aligned} \max_Z \quad & -Z_{11} \\ \text{s.t.} \quad & Z_{11} + 2Z_{23} = 1 \\ & Z_{22} = 0 \\ & Z \succeq 0. \end{aligned}$$

Since Z is positive semidefinite, then so must be the submatrix $\begin{pmatrix} Z_{22} & Z_{23} \\ Z_{23} & Z_{33} \end{pmatrix}$. But since $Z_{22} = 0$, then necessarily $Z_{23} = 0$. So, for any feasible point in the dual, the constraint $Z_{11} + Z_{23} = 1$ is equivalent to $Z_{11} = 1$. So, the dual attains an optimal point, but the value at optimality is different from the value at optimality of the primal problem.

■ L9.3 Conic solvers and software tools

Conic programming is a powerful tool that can be used to solve a wide range of optimization problems. Several excellent software packages are available that can solve conic programs in practice. Here are just a few:

- CVX [GB08] is a modeling language for convex optimization problems that can be used with several solvers, including MOSEK [AA00] and SCS [O'D+16]. It is compatible with Python and numpy.
- JuMP [Lub+23] is a modeling language for mathematical optimization problems in Julia. Like CVX, it can be used with several solvers, including MOSEK and SCS.
- MOSEK [AA00] is a commercial solver that can handle a wide range of optimization problems, including linear, quadratic, conic, and semidefinite programs.

■ L9.4 Further readings and bibliography

Hiriart-Urruty, J.-B., & Seeger, A. [HS10] wrote an excellent survey of the properties of the copositive cone.

For further reading on semidefinite programming, I find the short books by Klerk, E. de. [Kle02] and Gärtner, B., & Matoušek, J. [GM12] very well-written and approachable. The book by Boyd, S., & Vandenberghe, L. [BV04] is a great reference too.

The monograph by Ben-Tal, A., & Nemirovski, A. [BN01] contains an extensive list of applications of second-order cone problems and semidefinite problems.

- [Dia62] Diananda, P. H. (1962). On non-negative forms in real variables some or all of which are non-negative. *Mathematical Proceedings of the Cambridge Philosophical Society*, 58(1), 17–25. <https://doi.org/10.1017/S0305004100036185>
- [dP02] de Klerk, E., & Pasechnik, D. V. (2002). Approximating of the stability number of a graph via copositive programming. *SIAM J. Optim.*, 12(4), 875–892. <https://doi.org/10.1137/S105262340138324>
- [GM12] Gärtner, B., & Matoušek, J. (2012). *Approximation Algorithms and Semidefinite Programming*. Springer. <https://link.springer.com/book/10.1007/978-3-642-22015-9>
- [Kle02] Klerk, E. de. (2002). *Aspects of Semidefinite Programming*. Springer US. <https://link.springer.com/book/10.1007/b105286>

- [BV04] Boyd, S., & Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press. <https://web.stanford.edu/~boyd/cvxbook/>
- [GB08] Grant, M., & Boyd, S. (2008). *CVX: Matlab Software for Disciplined Convex Programming*.
- [AA00] Andersen, E. D., & Andersen, K. D. (2000). The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm. *High Performance Optimization*, 197–232.
- [O'D+16] O'Donoghue, B., Chu, E., Parikh, N., & Boyd, S. (2016). Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding. *Journal of Optimization Theory and Applications*, 169(3), 1042–1068.
- [Lub+23] Lubin, M., Dowson, O., Dias Garcia, J., Huchette, J., Legat, B., & Vielma, J. P. (2023). JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*. <https://doi.org/10.1007/s12532-023-00239-3>
- [HS10] Hiriart-Urruty, J.-B., & Seeger, A. (2010). A Variational Approach to Copositive Matrices. *SIAM Review*, 52(4), 593–629. <http://www.jstor.org/stable/41062014>
- [BN01] Ben-Tal, A., & Nemirovski, A. (2001). *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*. SIAM.

Changelog

- Mar 6, 2025: Fixed typo in conic dual problem (thanks Jonathan Huang!)
- Mar 9, 2025: Added remark about the hardness of even checking membership in the copositive cone.
- Mar 14, 2025: changed “determinant” into “principal minor” in Example L9.2 (<https://piazza.com/class/m6lg9aspoutda/post/95>).
- Mar 15, 2025: Fixed wrong sign in text at end of L9.2.4 (thanks Nicolas Gorlo!)
- Mar 19, 2025: Fixed typo “oder” → “order” (thanks Alina Yang!)

MIT OpenCourseWare
<https://ocw.mit.edu>

6.7220 Nonlinear Optimization
Spring 2025

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>