1

Convex Optimization Models: An Overview

Contents

In this chapter we provide an overview of some broad classes of convex optimization models. Our primary focus will be on large challenging problems, often connected in some way to duality. We will consider two types of duality. The first is *Lagrange duality* for constrained optimization, which is obtained by assigning dual variables to the constraints. The second is *Fenchel duality* together with its special case, conic duality, which involves a cost function that is the sum of two convex function components. Both of these duality structures arise often in applications, and in Sections 1.1 and 1.2 we provide an overview, and discuss some examples.†

In Sections 1.3 and 1.4, we discuss additional model structures involving a large number of additive terms in the cost, or a large number of constraints. These types of problems also arise often in the context of duality, as well as in other contexts such as machine learning and signal processing with large amounts of data. In Section 1.5, we discuss the exact penalty function technique, whereby we can transform a convex constrained optimization problem to an equivalent unconstrained problem.

1.1 LAGRANGE DUALITY

We start our overview of Lagrange duality with the basic case of nonlinear inequality constraints, and then consider extensions involving linear inequality and equality constraints. Consider the problem‡

$$\begin{array}{ll} \text{minimize} & f(x) \\ \text{subject to} & x \in X, \quad g(x) \leq 0, \end{array} \tag{1.1}$$

where X is a nonempty set,

$$g(x) = \big(g_1(x), \ldots, g_r(x)\big)',$$

and $f : X \mapsto \Re$ and $g_j : X \mapsto \Re$, $j = 1, \ldots, r$, are given functions. We refer to this as the *primal problem*, and we denote its optimal value by f^*. A vector x satisfying the constraints of the problem is referred to as *feasible*. The *dual* of problem (1.1) is given by

$$\begin{array}{ll} \text{maximize} & q(\mu) \\ \text{subject to} & \mu \in \Re^r, \end{array} \tag{1.2}$$

† Consistent with its overview character, this chapter contains few proofs, and refers frequently to the literature, and to Appendix B, which contains a full list of definitions and propositions (without proofs) relating to nonalgorithmic aspects of convex optimization. This list reflects and summarizes the content of the author's "Convex Optimization Theory" book [Ber09]. The proposition numbers of [Ber09] have been preserved, so all omitted proofs of propositions in Appendix B can be readily accessed from [Ber09].

‡ Appendix A contains an overview of the mathematical notation, terminology, and results from linear algebra and real analysis that we will be using.

where the dual function q is

$$q(\mu) = \begin{cases} \inf_{x \in X} L(x,\mu) & \text{if } \mu \geq 0, \\ -\infty & \text{otherwise,} \end{cases}$$

and L is the Lagrangian function defined by

$$L(x,\mu) = f(x) + \mu' g(x), \qquad x \in X, \ \mu \in \Re^r;$$

(cf. Section 5.3 of Appendix B).

Note that the dual function is extended real-valued, and that the effective constraint set of the dual problem is

$$\left\{ \mu \geq 0 \;\middle|\; \inf_{x \in X} L(x,\mu) > -\infty \right\}.$$

The optimal value of the dual problem is denoted by q^*.

The *weak duality* relation, $q^* \leq f^*$, always holds. It is easily shown by writing for all $\mu \geq 0$, and $x \in X$ with $g(x) \leq 0$,

$$q(\mu) = \inf_{z \in X} L(z,\mu) \leq L(x,\mu) = f(x) + \sum_{j=1}^{r} \mu_j g_j(x) \leq f(x),$$

so that

$$q^* = \sup_{\mu \in \Re^r} q(\mu) = \sup_{\mu \geq 0} q(\mu) \leq \inf_{x \in X,\, g(x) \leq 0} f(x) = f^*.$$

We state this formally as follows (cf. Prop. 4.1.2 in Appendix B).

Proposition 1.1.1: (Weak Duality Theorem) Consider problem (1.1). For any feasible solution x and any $\mu \in \Re^r$, we have $q(\mu) \leq f(x)$. Moreover, $q^* \leq f^*$.

When $q^* = f^*$, we say that *strong duality* holds. The following proposition gives necessary and sufficient conditions for strong duality, and primal and dual optimality (see Prop. 5.3.2 in Appendix B).

Proposition 1.1.2: (Optimality Conditions) Consider problem (1.1). There holds $q^* = f^*$, and (x^*, μ^*) are a primal and dual optimal solution pair if and only if x^* is feasible, $\mu^* \geq 0$, and

$$x^* \in \arg\min_{x \in X} L(x,\mu^*), \qquad \mu_j^* g_j(x^*) = 0, \quad j = 1, \ldots, r.$$

Both of the preceding propositions do not require any convexity assumptions on f, g, and X. However, generally the analytical and algorithmic solution process is simplified when strong duality ($q^* = f^*$) holds. This typically requires convexity assumptions, and in some cases conditions on $\mathrm{ri}(X)$, the relative interior of X, as exemplified by the following result, given in Prop. 5.3.1 in Appendix B. The result delineates the two principal cases where there is no duality gap in an inequality-constrained problem.

Proposition 1.1.3: (Strong Duality – Existence of Dual Optimal Solutions) Consider problem (1.1) under the assumption that the set X is convex, and the functions f, and $g_1, \ldots, g_r$ are convex. Assume further that f^* is finite, and that one of the following two conditions holds:

(1) There exists $\overline{x} \in X$ such that $g_j(\overline{x}) < 0$ for all $j = 1, \ldots, r$.

(2) The functions g_j, $j = 1, \ldots, r$, are affine, and there exists $\overline{x} \in \mathrm{ri}(X)$ such that $g(\overline{x}) \leq 0$.

Then $q^* = f^*$ and there exists at least one dual optimal solution. Under condition (1) the set of dual optimal solutions is also compact.

Convex Programming with Inequality and Equality Constraints

Let us consider an extension of problem (1.1), with additional linear equality constraints. It is our principal constrained optimization model under convexity assumptions, and it will be referred to as the *convex programming problem*. It is given by

$$\begin{aligned} &\text{minimize} \quad f(x) \\ &\text{subject to} \quad x \in X, \quad g(x) \leq 0, \quad Ax = b, \end{aligned} \tag{1.3}$$

where X is a convex set, $g(x) = \big(g_1(x), \ldots, g_r(x)\big)'$, $f : X \mapsto \Re$ and $g_j : X \mapsto \Re$, $j = 1, \ldots, r$, are given convex functions, A is an $m \times n$ matrix, and $b \in \Re^m$.

The preceding duality framework may be applied to this problem by converting the constraint $Ax = b$ to the equivalent set of linear inequality constraints

$$Ax \leq b, \qquad -Ax \leq -b,$$

with corresponding dual variables $\lambda^+ \geq 0$ and $\lambda^- \geq 0$. The Lagrangian function is

$$f(x) + \mu' g(x) + (\lambda^+ - \lambda^-)'(Ax - b),$$

and by introducing a dual variable

$$\lambda = \lambda^+ - \lambda^-$$

with no sign restriction, it can be written as

$$L(x, \mu, \lambda) = f(x) + \mu' g(x) + \lambda'(Ax - b).$$

The dual problem is

$$\begin{array}{ll} \text{maximize} & \inf_{x \in X} L(x, \mu, \lambda) \\ \text{subject to} & \mu \geq 0, \ \lambda \in \Re^m. \end{array}$$

In this manner, Prop. 1.1.3 under condition (2), together with Prop. 1.1.2, yield the following for the case where all constraint functions are linear.

Proposition 1.1.4: (Convex Programming – Linear Equality and Inequality Constraints) Consider problem (1.3).

(a) Assume that f^* is finite, that the functions g_j are affine, and that there exists $\overline{x} \in \mathrm{ri}(X)$ such that $A\overline{x} = b$ and $g(\overline{x}) \leq 0$. Then $q^* = f^*$ and there exists at least one dual optimal solution.

(b) There holds $f^* = q^*$, and (x^*, μ^*, λ^*) are a primal and dual optimal solution pair if and only if x^* is feasible, $\mu^* \geq 0$, and

$$x^* \in \arg\min_{x \in X} L(x, \mu^*, \lambda^*), \qquad \mu_j^* g_j(x^*) = 0, \quad j = 1, \ldots, r.$$

In the special case where there are no inequality constraints:

$$\begin{array}{ll} \text{minimize} & f(x) \\ \text{subject to} & x \in X, \quad Ax = b, \end{array} \tag{1.4}$$

the Lagrangian function is

$$L(x, \lambda) = f(x) + \lambda'(Ax - b),$$

and the dual problem is

$$\begin{array}{ll} \text{maximize} & \inf_{x \in X} L(x, \lambda) \\ \text{subject to} & \lambda \in \Re^m. \end{array}$$

The corresponding result, a simpler special case of Prop. 1.1.4, is given in the following proposition.

Proposition 1.1.5: (Convex Programming – Linear Equality Constraints) Consider problem (1.4).

(a) Assume that f^* is finite and that there exists $\overline{x} \in \mathrm{ri}(X)$ such that $A\overline{x} = b$. Then $f^* = q^*$ and there exists at least one dual optimal solution.

(b) There holds $f^* = q^*$, and (x^*, λ^*) are a primal and dual optimal solution pair if and only if x^* is feasible and

$$x^* \in \arg\min_{x \in X} L(x, \lambda^*).$$

The following is an extension of Prop. 1.1.4(a) to the case where the inequality constraints may be nonlinear. It is the most general convex programming result relating to duality in this section (see Prop. 5.3.5 in Appendix B).

Proposition 1.1.6: (Convex Programming – Linear Equality and Nonlinear Inequality Constraints) Consider problem (1.3). Assume that f^* is finite, that there exists $\overline{x} \in X$ such that $A\overline{x} = b$ and $g(\overline{x}) < 0$, and that there exists $\tilde{x} \in \mathrm{ri}(X)$ such that $A\tilde{x} = b$. Then $q^* = f^*$ and there exists at least one dual optimal solution.

Aside from the preceding results, there are alternative optimality conditions for convex and nonconvex optimization problems, which are based on extended versions of the Fritz John theorem; see [BeO02] and [BOT06], and the textbooks [Ber99] and [BNO03]. These conditions are derived using a somewhat different line of analysis and supplement the ones given here, but we will not have occasion to use them in this book.

Discrete Optimization and Lower Bounds

The preceding propositions deal mostly with situations where strong duality holds ($q^* = f^*$). However, duality can be useful even when there is duality gap, as often occurs in problems that have a finite constraint set X. An example is *integer programming*, where the components of x must be integers from a bounded range (usually 0 or 1). An important special case is the linear 0-1 integer programming problem

$$\begin{aligned} &\text{minimize} \quad c'x \\ &\text{subject to} \quad Ax \le b, \quad x_i = 0 \text{ or } 1, \quad i = 1, \ldots, n, \end{aligned}$$

where $x = (x_1, \ldots, x_n)$.

A principal approach for solving discrete optimization problems with a finite constraint set is the *branch-and-bound method*, which is described in many sources; see e.g., one of the original works [LaD60], the survey [BaT85], and the book [NeW88]. The general idea of the method is that bounds on the cost function can be used to exclude from consideration portions of the feasible set. To illustrate, consider minimizing $F(x)$ over $x \in X$, and let Y_1, Y_2 be two subsets of X. Suppose that we have bounds

$$\underline{F}_1 \leq \min_{x \in Y_1} f(x), \qquad \overline{F}_2 \geq \min_{x \in Y_2} f(x).$$

Then, if $\overline{F}_2 \leq \underline{F}_1$, the solutions in Y_1 may be disregarded since their cost cannot be smaller than the cost of the best solution in Y_2. The lower bound $\underline{F}_1$ can often be conveniently obtained by minimizing f over a suitably enlarged version of Y_1, while for the upper bound $\overline{F}_2$, a value $f(x)$, where $x \in Y_2$, may be used.

Branch-and-bound is often based on weak duality (cf. Prop. 1.1.1) to obtain lower bounds to the optimal cost of restricted problems of the form

$$\begin{aligned} &\text{minimize} \quad f(x) \\ &\text{subject to} \quad x \in \tilde{X}, \qquad g(x) \leq 0, \end{aligned} \tag{1.5}$$

where $\tilde{X}$ is a subset of X; for example in the 0-1 integer case where X specifies that all x_i should be 0 or 1, $\tilde{X}$ may be the set of all 0-1 vectors x such that one or more components x_i are fixed at either 0 or 1 (i.e., are restricted to satisfy $x_i = 0$ for all $x \in \tilde{X}$ or $x_i = 1$ for all $x \in \tilde{X}$). These lower bounds can often be obtained by finding a dual-feasible (possibly dual-optimal) solution $\mu \geq 0$ of this problem and the corresponding dual value

$$q(\mu) = \inf_{x \in \tilde{X}} \big\{ f(x) + \mu' g(x) \big\}, \tag{1.6}$$

which by weak duality, is a lower bound to the optimal value of the restricted problem (1.5). In a strengthened version of this approach, the given inequality constraints $g(x) \leq 0$ may be augmented by additional inequalities that are known to be satisfied by optimal solutions of the original problem.

An important point here is that when $\tilde{X}$ is finite, the dual function q of Eq. (1.6) is concave and polyhedral. Thus solving the dual problem amounts to minimizing the polyhedral function $-q$ over the nonnegative orthant. This is a major context within which polyhedral functions arise in convex optimization.

1.1.1 Separable Problems – Decomposition

Let us now discuss an important problem structure that involves Lagrange duality and arises frequently in applications. Here x has m components,

$x = (x_1, \ldots, x_m)$, with each x_i being a vector of dimension n_i (often $n_i = 1$). The problem has the form

$$\begin{aligned} &\text{minimize} \quad \sum_{i=1}^{m} f_i(x_i) \\ &\text{subject to} \quad \sum_{i=1}^{m} g_{ij}(x_i) \le 0, \quad x_i \in X_i, \quad i = 1, \ldots, m, \ j = 1, \ldots, r, \end{aligned} \tag{1.7}$$

where $f_i : \Re^{n_i} \mapsto \Re$ and $g_{ij} : \Re^{n_i} \mapsto \Re^r$ are given functions, and X_i are given subsets of $\Re^{n_i}$. By assigning a dual variable μ_j to the jth constraint, we obtain the dual problem [cf. Eq. (1.2)]

$$\begin{aligned} &\text{maximize} \quad \sum_{i=1}^{m} q_i(\mu) \\ &\text{subject to} \quad \mu \ge 0, \end{aligned} \tag{1.8}$$

where

$$q_i(\mu) = \inf_{x_i \in X_i} \left\{ f_i(x_i) + \sum_{j=1}^{r} \mu_j g_{ij}(x_i) \right\},$$

and $\mu = (\mu_1, \ldots, \mu_r)$.

Note that the minimization involved in the calculation of the dual function has been decomposed into m simpler minimizations. These minimizations are often conveniently done either analytically or computationally, in which case the dual function can be easily evaluated. This is the key advantageous structure of separable problems: it facilitates computation of dual function values (as well as subgradients as we will see in Section 3.1), and it is amenable to decomposition and distributed computation.

Let us also note that in the special case where the components x_i are one-dimensional, and the functions f_i and sets X_i are convex, there is a particularly favorable duality result for the separable problem (1.7): essentially, strong duality holds without any qualifications such as the linearity of the constraint functions, or the Slater condition of Prop. 1.1.3; see [Tse09].

Duality Gap Estimates for Nonconvex Separable Problems

The separable structure is additionally helpful when the cost and/or the constraints are not convex, and there is a duality gap. In particular, in this case *the duality gap turns out to be relatively small and can often be shown to diminish to zero relative to the optimal primal value as the number m of separable terms increases*. As a result, one can often obtain a near-optimal primal solution, starting from a dual-optimal solution, without resorting to costly branch-and-bound procedures.

The small duality gap size is a consequence of the structure of the set S of constraint-cost pairs of problem (1.7), which in the case of a separable problem, can be written as a vector sum of m sets, one for each separable term, i.e.,

$$S = S_1 + \cdots + S_m,$$

where

$$S_i = \big\{ \big(g_i(x_i), f_i(x_i)\big) \mid x_i \in X_i \big\},$$

and $g_i : \Re^{n_i} \mapsto \Re^r$ is the function $g_i(x_i) = \big(g_{i1}(x_i), \ldots, g_{im}(x_i)\big)$. It can be shown that the duality gap is related to how much S "differs" from its convex hull (a geometric explanation is given in [Ber99], Section 5.1.6, and [Ber09], Section 5.7). Generally, a set that is the vector sum of a large number of possibly nonconvex but roughly similar sets "tends to be convex" in the sense that any vector in its convex hull can be closely approximated by a vector in the set. As a result, the duality gap tends to be relatively small. The analytical substantiation is based on a theorem by Shapley and Folkman (see [Ber99], Section 5.1, or [Ber09], Prop. 5.7.1, for a statement and proof of this theorem). In particular, it is shown in [AuE76], and also [BeS82], [Ber82a], Section 5.6.1, under various reasonable assumptions, that the duality gap satisfies

$$f^* - q^* \leq (r+1) \max_{i=1,\ldots,m} \rho_i,$$

where for each i, ρ_i is a nonnegative scalar that depends on the structure of the functions f_i, g_{ij}, $j = 1, \ldots, r$, and the set X_i (the paper [AuE76] focuses on the case where the problem is nonconvex but continuous, while [BeS82] and [Ber82a] focus on an important class of mixed integer programming problems). This estimate suggests that as $m \to \infty$ and $|f^*| \to \infty$, the duality gap is bounded, while the "relative" duality gap $(f^* - q^*)/|f^*|$ diminishes to 0 as $m \to \infty$.

The duality gap has also been investigated in the author's book [Ber09] within the more general min common-max crossing framework (Section 4.1 of Appendix B). This framework includes as special cases minimax and zero-sum game problems. In particular, consider a function $\phi : X \times Z \mapsto \Re$ defined over nonempty subsets $X \subset \Re^n$ and $Z \subset \Re^m$. Then it can be shown that the gap between "infsup" and "supinf" of ϕ can be decomposed into the sum of two terms that can be computed separately: one term can be attributed to the lack of convexity and/or closure of ϕ with respect to x, and the other can be attributed to the lack of concavity and/or upper semicontinuity of ϕ with respect to z. We refer to [Ber09], Section 5.7.2, for the analysis.

1.1.2 Partitioning

It is important to note that there are several different ways to introduce duality in the solution of large-scale optimization problems. For example a strategy, often called *partitioning*, is to divide the variables in two subsets, and minimize first with respect to one subset while taking advantage of whatever simplification may arise by fixing the variables in the other subset.

As an example, the problem

$$\begin{aligned} &\text{minimize} \quad F(x) + G(y) \\ &\text{subject to} \quad Ax + By = c, \quad x \in X, \quad y \in Y, \end{aligned}$$

can be written as

$$\begin{aligned} &\text{minimize} \quad F(x) + \inf_{By = c - Ax,\, y \in Y} G(y) \\ &\text{subject to} \quad x \in X, \end{aligned}$$

or

$$\begin{aligned} &\text{minimize} \quad F(x) + p(c - Ax) \\ &\text{subject to} \quad x \in X, \end{aligned}$$

where p is given by

$$p(u) = \inf_{By = u,\, y \in Y} G(y).$$

In favorable cases, p can be dealt with conveniently (see e.g., the book [Las70] and the paper [Geo72]).

Strategies of splitting or transforming the variables to facilitate algorithmic solution will be frequently encountered in what follows, and in a variety of contexts, including duality. The next section describes some significant contexts of this type.

1.2 FENCHEL DUALITY AND CONIC PROGRAMMING

Let us consider the Fenchel duality framework (see Section 5.3.5 of Appendix B). It involves the problem

$$\begin{aligned} &\text{minimize} \quad f_1(x) + f_2(Ax) \\ &\text{subject to} \quad x \in \Re^n, \end{aligned} \tag{1.9}$$

where A is an $m \times n$ matrix, $f_1 : \Re^n \mapsto (-\infty, \infty]$ and $f_2 : \Re^m \mapsto (-\infty, \infty]$ are closed proper convex functions, and we assume that there exists a feasible solution, i.e., an $x \in \Re^n$ such that $x \in \text{dom}(f_1)$ and $Ax \in \text{dom}(f_2)$.†

The problem is equivalent to the following constrained optimization problem in the variables $x_1 \in \Re^n$ and $x_2 \in \Re^m$:

$$\begin{aligned} &\text{minimize} \quad f_1(x_1) + f_2(x_2) \\ &\text{subject to} \quad x_1 \in \text{dom}(f_1), \quad x_2 \in \text{dom}(f_2), \qquad x_2 = Ax_1. \end{aligned} \tag{1.10}$$

† We remind the reader that our convex analysis notation, terminology, and nonalgorithmic theory are summarized in Appendix B.

Viewing this as a convex programming problem with the linear equality constraint $x_2 = Ax_1$, we obtain the dual function as

$$\begin{aligned} q(\lambda) &= \inf_{x_1 \in \mathrm{dom}(f_1),\ x_2 \in \mathrm{dom}(f_2)} \big\{ f_1(x_1) + f_2(x_2) + \lambda'(x_2 - Ax_1) \big\} \\ &= \inf_{x_1 \in \Re^n} \big\{ f_1(x_1) - \lambda' A x_1 \big\} + \inf_{x_2 \in \Re^m} \big\{ f_2(x_2) + \lambda' x_2 \big\}. \end{aligned}$$

The dual problem of maximizing q over $\lambda \in \Re^m$, after a sign change to convert it to a minimization problem, takes the form

$$\begin{aligned} &\text{minimize} \quad f_1^\star(A'\lambda) + f_2^\star(-\lambda) \\ &\text{subject to} \quad \lambda \in \Re^m, \end{aligned} \tag{1.11}$$

where $f_1^\star$ and $f_2^\star$ are the conjugate functions of f_1 and f_2. We denote by f^* and q^* the corresponding optimal primal and dual values [q^* is the negative of the optimal value of problem (1.11)].

The following Fenchel duality result is given as Prop. 5.3.8 in Appendix B. Parts (a) and (b) are obtained by applying Prop. 1.1.5(a) to problem (1.10), viewed as a problem with $x_2 = Ax_1$ as the only linear equality constraint. The first equation of part (c) is a consequence of Prop. 1.1.5(b). Its equivalence with the last two equations is a consequence of the Conjugate Subgradient Theorem (Prop. 5.4.3, App. B), which states that for a closed proper convex function f, its conjugate $f^\star$, and any pair of vectors (x, y), we have

$$x \in \arg\min_{z \in \Re^n} \big\{ f(z) - z'y \big\} \quad \text{iff} \quad y \in \partial f(x) \quad \text{iff} \quad x \in \partial f^\star(y),$$

with all of these three relations being equivalent to $x'y = f(x) + f^\star(y)$. Here $\partial f(x)$ denotes the subdifferential of f at x (the set of all subgradients of f at x); see Section 5.4 of Appendix B.

Proposition 1.2.1: (Fenchel Duality) Consider problem (1.9).

(a) If f^* is finite and $\big(A \cdot \mathrm{ri}\big(\mathrm{dom}(f_1)\big)\big) \cap \mathrm{ri}\big(\mathrm{dom}(f_2)\big) \neq \emptyset$, then $f^* = q^*$ and there exists at least one dual optimal solution.

(b) If q^* is finite and $\mathrm{ri}\big(\mathrm{dom}(f_1^\star)\big) \cap \big(A' \cdot \mathrm{ri}\big(-\mathrm{dom}(f_2^\star)\big)\big) \neq \emptyset$, then $f^* = q^*$ and there exists at least one primal optimal solution.

(c) There holds $f^* = q^*$, and (x^*, λ^*) is a primal and dual optimal solution pair if and only if any one of the following three equivalent conditions hold:

$$x^* \in \arg\min_{x \in \Re^n} \big\{ f_1(x) - x'A'\lambda^* \big\} \quad \text{and} \quad Ax^* \in \arg\min_{z \in \Re^m} \big\{ f_2(z) + z'\lambda^* \big\}, \tag{1.12}$$

$$A'\lambda^* \in \partial f_1(x^*) \quad \text{and} \quad -\lambda^* \in \partial f_2(Ax^*), \tag{1.13}$$

$$x^* \in \partial f_1^\star(A'\lambda^*) \quad \text{and} \quad Ax^* \in \partial f_2^\star(-\lambda^*). \tag{1.14}$$

Minimax Problems

Minimax problems involve minimization over a set X of a function $\overline{F}$ of the form

$$\overline{F}(x) = \sup_{z \in Z} \phi(x,z),$$

where X and Z are subsets of $\Re^n$ and $\Re^m$, respectively, and $\phi : \Re^n \times \Re^m \mapsto \Re$ is a given function. Some (but not all) problems of this type are related to constrained optimization and Fenchel duality.

Example 1.2.1: (Connection with Constrained Optimization)

Let ϕ and Z have the form

$$\phi(x,z) = f(x) + z'g(x), \qquad Z = \{z \mid z \geq 0\},$$

where $f : \Re^n \mapsto \Re$ and $g : \Re^n \mapsto \Re^m$ are given functions. Then it is seen that

$$\overline{F}(x) = \sup_{z \in Z} \phi(x,z) = \begin{cases} f(x) & \text{if } g(x) \leq 0, \\ \infty & \text{otherwise.} \end{cases}$$

Thus minimization of $\overline{F}$ over $x \in X$ is equivalent to solving the constrained optimization problem

$$\begin{aligned} &\text{minimize} \quad f(x) \\ &\text{subject to} \quad x \in X, \quad g(x) \leq 0. \end{aligned} \tag{1.15}$$

The dual problem is to maximize over $z \geq 0$ the function

$$\underline{F}(z) = \inf_{x \in X} \big\{ f(x) + z'g(x) \big\} = \inf_{x \in X} \phi(x,z),$$

and the minimax equality

$$\inf_{x \in X} \sup_{z \in Z} \phi(x,z) = \sup_{z \in Z} \inf_{x \in X} \phi(x,z) \tag{1.16}$$

is equivalent to problem (1.15) having no duality gap.

Example 1.2.2: (Connection with Fenchel Duality)

Let ϕ have the special form

$$\phi(x,z) = f(x) + z'Ax - g(z),$$

where $f : \Re^n \mapsto \Re$ and $g : \Re^m \mapsto \Re$ are given functions, and A is a given $m \times n$ matrix. Then we have

$$\overline{F}(x) = \sup_{z \in Z} \phi(x,z) = f(x) + \sup_{z \in Z} \big\{ (Ax)'z - g(z) \big\} = f(x) + \hat{g}^\star(Ax),$$

where $\hat{g}^\star$ is the conjugate of the function

$$\hat{g}(z) = \begin{cases} g(z) & \text{if } z \in Z, \\ \infty & \text{otherwise.} \end{cases}$$

Thus the minimax problem of minimizing $\overline{F}$ over $x \in X$ comes under the Fenchel framework (1.9) with $f_2 = \hat{g}^\star$ and f_1 given by

$$f_1(x) = \begin{cases} f(x) & \text{if } x \in X, \\ \infty & \text{if } x \notin X. \end{cases}$$

It can also be verified that the Fenchel dual problem (1.11) is equivalent to maximizing over $z \in Z$ the function $\underline{F}(z) = \inf_{x \in X} \phi(x, z)$. Again having no duality gap is equivalent to the minimax equality (1.16) holding.

Finally note that strong duality theory is connected with minimax problems primarily when X and Z are convex sets, and ϕ is convex in x and concave in z. When Z is a finite set, there is a different connection with constrained optimization that does not involve Fenchel duality and applies without any convexity conditions. In particular, the problem

$$\begin{aligned} &\text{minimize} \quad \max\big\{g_1(x), \ldots, g_r(x)\big\} \\ &\text{subject to} \quad x \in X, \end{aligned}$$

where $g_j : \Re^n \mapsto \Re$ are any real-valued functions, is equivalent to the constrained optimization problem

$$\begin{aligned} &\text{minimize} \quad y \\ &\text{subject to} \quad x \in X, \qquad g_j(x) \le y, \quad j = 1, \ldots, r, \end{aligned}$$

where y is an additional scalar optimization variable. Minimax problems will be discussed further later, in Section 1.4, as an example of problems that may involve a large number of constraints.

Conic Programming

An important problem structure, which can be analyzed as a special case of the Fenchel duality framework is *conic programming*. This is the problem

$$\begin{aligned} &\text{minimize} \quad f(x) \\ &\text{subject to} \quad x \in C, \end{aligned} \tag{1.17}$$

where $f : \Re^n \mapsto (-\infty, \infty]$ is a closed proper convex function and C is a closed convex cone in $\Re^n$.

Indeed, let us apply Fenchel duality with A equal to the identity and the definitions

$$f_1(x) = f(x), \qquad f_2(x) = \begin{cases} 0 & \text{if } x \in C, \\ \infty & \text{if } x \notin C. \end{cases}$$

The corresponding conjugates are

$$f_1^\star(\lambda) = \sup_{x\in\Re^n}\big\{\lambda'x - f(x)\big\}, \qquad f_2^\star(\lambda) = \sup_{x\in C}\lambda'x = \begin{cases} 0 & \text{if } \lambda\in C^*,\\ \infty & \text{if } \lambda\notin C^*,\end{cases}$$

where

$$C^* = \{\lambda \mid \lambda'x \le 0,\ \forall\ x\in C\}$$

is the polar cone of C (note that $f_2^\star$ is the support function of C; cf. Section 1.6 of Appendix B). The dual problem is

$$\begin{aligned} &\text{minimize} && f^\star(\lambda)\\ &\text{subject to} && \lambda\in\hat{C}, \end{aligned} \tag{1.18}$$

where $f^\star$ is the conjugate of f and $\hat{C}$ is the negative polar cone (also called the *dual cone* of C):

$$\hat{C} = -C^* = \{\lambda \mid \lambda'x \ge 0,\ \forall\ x\in C\}.$$

Note the symmetry between primal and dual problems. The strong duality relation $f^* = q^*$ can be written as

$$\inf_{x\in C} f(x) = -\inf_{\lambda\in\hat{C}} f^\star(\lambda).$$

The following proposition translates the conditions of Prop. 1.2.1(a), which guarantees that there is no duality gap and that the dual problem has an optimal solution.

Proposition 1.2.2: (Conic Duality Theorem) Assume that the primal conic problem (1.17) has finite optimal value, and moreover $\mathrm{ri}\big(\mathrm{dom}(f)\big)\cap\mathrm{ri}(C)\neq\emptyset$. Then, there is no duality gap and the dual problem (1.18) has an optimal solution.

Using the symmetry of the primal and dual problems, we also obtain that there is no duality gap and the primal problem (1.17) has an optimal solution if the optimal value of the dual conic problem (1.18) is finite and $\mathrm{ri}\big(\mathrm{dom}(f^\star)\big)\cap\mathrm{ri}(\hat{C})\neq\emptyset$. It is also possible to derive primal and dual optimality conditions by translating the optimality conditions of the Fenchel duality framework [Prop. 1.2.1(c)].

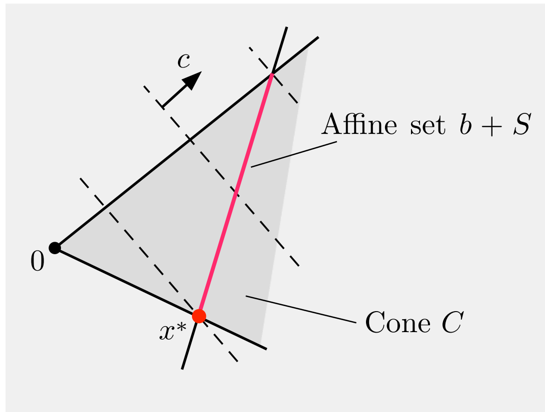

Figure 1.2.1. Illustration of a linear-conic problem: minimizing a linear function $c'x$ over the intersection of an affine set $b + S$ and a convex cone C.

1.2.1 Linear-Conic Problems

An important special case of conic programming, called *linear-conic problem*, arises when $\mathrm{dom}(f)$ is an affine set and f is linear over $\mathrm{dom}(f)$, i.e.,

$$f(x) = \begin{cases} c'x & \text{if } x \in b + S, \\ \infty & \text{if } x \notin b + S, \end{cases}$$

where b and c are given vectors, and S is a subspace. Then the primal problem can be written as

$$\begin{aligned} &\text{minimize} \quad c'x \\ &\text{subject to} \quad x - b \in S, \quad x \in C; \end{aligned} \tag{1.19}$$

see Fig. 1.2.1.

To derive the dual problem, we note that

$$\begin{aligned} f^\star(\lambda) &= \sup_{x-b\in S} (\lambda - c)'x \\ &= \sup_{y\in S} (\lambda - c)'(y + b) \\ &= \begin{cases} (\lambda - c)'b & \text{if } \lambda - c \in S^\perp, \\ \infty & \text{if } \lambda - c \notin S^\perp. \end{cases} \end{aligned}$$

It can be seen that the dual problem $\min_{\lambda\in\hat{C}} f^\star(\lambda)$ [cf. Eq. (1.18)], after discarding the superfluous term $c'b$ from the cost, can be written as

$$\begin{aligned} &\text{minimize} \quad b'\lambda \\ &\text{subject to} \quad \lambda - c \in S^\perp, \quad \lambda \in \hat{C}, \end{aligned} \tag{1.20}$$

where $\hat{C}$ is the dual cone:

$$\hat{C} = \{\lambda \mid \lambda' x \geq 0,\ \forall\ x \in C\}.$$

By specializing the conditions of the Conic Duality Theorem (Prop. 1.2.2) to the linear-conic duality context, we obtain the following.

Proposition 1.2.3: (Linear-Conic Duality Theorem) Assume that the primal problem (1.19) has finite optimal value, and moreover $(b+S)\cap \mathrm{ri}(C) \neq \emptyset$. Then, there is no duality gap and the dual problem has an optimal solution.

Special Forms of Linear-Conic Problems

The primal and dual linear-conic problems (1.19) and (1.20) have been placed in an elegant symmetric form. There are also other useful formats that parallel and generalize similar formats in linear programming. For example, we have the following dual problem pairs:

$$\min_{Ax=b,\ x\in C} c'x \qquad \Longleftrightarrow \qquad \max_{c-A'\lambda\in \hat{C}} b'\lambda, \tag{1.21}$$

$$\min_{Ax-b\in C} c'x \qquad \Longleftrightarrow \qquad \max_{A'\lambda=c,\ \lambda\in \hat{C}} b'\lambda, \tag{1.22}$$

where A is an $m \times n$ matrix, and $x \in \Re^n$, $\lambda \in \Re^m$, $c \in \Re^n$, $b \in \Re^m$.

To verify the duality relation (1.21), let $\overline{x}$ be any vector such that $A\overline{x} = b$, and let us write the primal problem on the left in the primal conic form (1.19) as

$$\begin{aligned} &\text{minimize} \quad c'x \\ &\text{subject to} \quad x - \overline{x} \in \mathrm{N}(A), \quad x \in C, \end{aligned}$$

where $\mathrm{N}(A)$ is the nullspace of A. The corresponding dual conic problem (1.20) is to solve for μ the problem

$$\begin{aligned} &\text{minimize} \quad \overline{x}'\mu \\ &\text{subject to} \quad \mu - c \in \mathrm{N}(A)^\perp, \quad \mu \in \hat{C}. \end{aligned} \tag{1.23}$$

Since $\mathrm{N}(A)^\perp$ is equal to $\mathrm{Ra}(A')$, the range of A', the constraints of problem (1.23) can be equivalently written as $c - \mu \in -\mathrm{Ra}(A') = \mathrm{Ra}(A')$, $\mu \in \hat{C}$, or

$$c - \mu = A'\lambda, \qquad \mu \in \hat{C},$$

for some $\lambda \in \Re^m$. Making the change of variables $\mu = c - A'\lambda$, the dual problem (1.23) can be written as

$$\begin{aligned} &\text{minimize} \quad \overline{x}'(c - A'\lambda) \\ &\text{subject to} \quad c - A'\lambda \in \hat{C}. \end{aligned}$$

By discarding the constant $\overline{x}'c$ from the cost function, using the fact $A\overline{x} = b$, and changing from minimization to maximization, we see that this dual problem is equivalent to the one in the right-hand side of the duality pair (1.21). The duality relation (1.22) is proved similarly.

We next discuss two important special cases of conic programming: *second order cone programming* and *semidefinite programming*. These problems involve two different special cones, and an explicit definition of the affine set constraint. They arise in a variety of applications, and their computational difficulty in practice tends to lie between that of linear and quadratic programming on one hand, and general convex programming on the other hand.

1.2.2 Second Order Cone Programming

In this section we consider the linear-conic problem (1.22), with the cone

$$C = \left\{ (x_1, \ldots, x_n) \;\middle|\; x_n \geq \sqrt{x_1^2 + \cdots + x_{n-1}^2} \right\},$$

which is known as the *second order cone* (see Fig. 1.2.2). The dual cone is

$$\hat{C} = \{ y \mid 0 \leq y'x,\ \forall\ x \in C \} = \left\{ y \;\middle|\; 0 \leq \inf_{\|(x_1,\ldots,x_{n-1})\| \leq x_n} y'x \right\},$$

and it can be shown that $\hat{C} = C$. This property is referred to as *self-duality* of the second order cone, and is fairly evident from Fig. 1.2.2. For a proof, we write

$$\begin{aligned} \inf_{\|(x_1,\ldots,x_{n-1})\| \leq x_n} y'x &= \inf_{x_n \geq 0} \left\{ y_n x_n + \inf_{\|(x_1,\ldots,x_{n-1})\| \leq x_n} \sum_{i=1}^{n-1} y_i x_i \right\} \\ &= \inf_{x_n \geq 0} \big\{ y_n x_n - \|(y_1, \ldots, y_{n-1})\|\, x_n \big\} \\ &= \begin{cases} 0 & \text{if } \|(y_1, \ldots, y_{n-1})\| \leq y_n, \\ -\infty & \text{otherwise,} \end{cases} \end{aligned}$$

where the second equality follows because the minimum of the inner product of a vector $z \in \Re^{n-1}$ with vectors in the unit ball of $\Re^{n-1}$ is $-\|z\|$. Combining the preceding two relations, we have

$$y \in \hat{C} \quad \text{if and only if} \quad 0 \leq y_n - \|(y_1, \ldots, y_{n-1})\|,$$

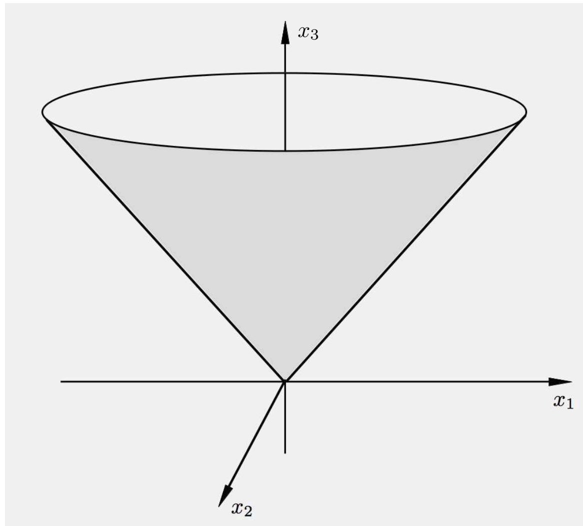

Figure 1.2.2. The second order cone

$$C = \left\{ (x_1, \ldots, x_n) \;\middle|\; x_n \geq \sqrt{x_1^2 + \cdots + x_{n-1}^2} \right\},$$

in $\Re^3$.

so $\hat{C} = C$.

The second order cone programming problem (SOCP for short) is

$$\begin{aligned} &\text{minimize} \quad c'x \\ &\text{subject to} \quad A_i x - b_i \in C_i, \ \ i = 1, \ldots, m, \end{aligned} \tag{1.24}$$

where $x \in \Re^n$, c is a vector in $\Re^n$, and for $i = 1, \ldots, m$, A_i is an $n_i \times n$ matrix, b_i is a vector in $\Re^{n_i}$, and C_i is the second order cone of $\Re^{n_i}$. It is seen to be a special case of the primal problem in the left-hand side of the duality relation (1.22), where

$$A = \begin{pmatrix} A_1 \\ \vdots \\ A_m \end{pmatrix}, \qquad b = \begin{pmatrix} b_1 \\ \vdots \\ b_m \end{pmatrix}, \qquad C = C_1 \times \cdots \times C_m.$$

Note that linear inequality constraints of the form $a_i'x - b_i \geq 0$ can be written as

$$\begin{pmatrix} 0 \\ a_i' \end{pmatrix} x - \begin{pmatrix} 0 \\ b_i \end{pmatrix} \in C_i,$$

where C_i is the second order cone of $\Re^2$. As a result, linear-conic problems involving second order cones contain as special cases linear programming problems.

We now observe that from the right-hand side of the duality relation (1.22), and the self-duality relation $C = \hat{C}$, the corresponding dual linear-conic problem has the form

$$\begin{aligned} &\text{maximize} \quad \sum_{i=1}^m b_i'\lambda_i \\ &\text{subject to} \quad \sum_{i=1}^m A_i'\lambda_i = c, \quad \lambda_i \in C_i,\ i = 1,\ldots,m, \end{aligned} \tag{1.25}$$

where $\lambda = (\lambda_1,\ldots,\lambda_m)$. By applying the Linear-Conic Duality Theorem (Prop. 1.2.3), we have the following.

Proposition 1.2.4: (Second Order Cone Duality Theorem) Consider the primal SOCP (1.24), and its dual problem (1.25).

(a) If the optimal value of the primal problem is finite and there exists a feasible solution $\overline{x}$ such that

$$A_i\overline{x} - b_i \in \text{int}(C_i), \qquad i = 1,\ldots,m,$$

then there is no duality gap, and the dual problem has an optimal solution.

(b) If the optimal value of the dual problem is finite and there exists a feasible solution $\overline{\lambda} = (\overline{\lambda}_1,\ldots,\overline{\lambda}_m)$ such that

$$\overline{\lambda}_i \in \text{int}(C_i), \qquad i = 1,\ldots,m,$$

then there is no duality gap, and the primal problem has an optimal solution.

Note that while the Linear-Conic Duality Theorem requires a relative interior point condition, the preceding proposition requires an interior point condition. The reason is that the second order cone has nonempty interior, so its relative interior coincides with its interior.

The SOCP arises in many application contexts, and significantly, it can be solved numerically with powerful specialized algorithms that belong to the class of interior point methods, which will be discussed in Section 6.8. We refer to the literature for a more detailed description and analysis (see e.g., the books [BeN01], [BoV04]).

Generally, SOCPs can be recognized from the presence of convex quadratic functions in the cost or the constraint functions. The following are illustrative examples. The first example relates to the field of robust optimization, which involves optimization under uncertainty described by set membership.

Example 1.2.3: (Robust Linear Programming)

Frequently, there is uncertainty about the data of an optimization problem, so one would like to have a solution that is adequate for a whole range of the uncertainty. A popular formulation of this type, is to assume that the constraints contain parameters that take values in a given set, and require that the constraints are satisfied for all values in that set. This approach is also known as a set membership description of the uncertainty and has been used in fields other than optimization, such as set membership estimation, and minimax control (see the textbook [Ber07], which also surveys earlier work).

As an example, consider the problem

$$\begin{aligned} &\text{minimize} \ \ c'x \\ &\text{subject to} \ \ a_j'x \le b_j, \quad \forall\ (a_j,b_j) \in T_j, \quad j=1,\ldots,r, \end{aligned} \tag{1.26}$$

where $c \in \Re^n$ is a given vector, and T_j is a given subset of $\Re^{n+1}$ to which the constraint parameter vectors (a_j, b_j) must belong. The vector x must be chosen so that the constraint $a_j'x \le b_j$ is satisfied for all $(a_j, b_j) \in T_j$, $j = 1,\ldots,r$.

Generally, when T_j contains an infinite number of elements, this problem involves a correspondingly infinite number of constraints. To convert the problem to one involving a finite number of constraints, we note that

$$a_j'x \le b_j, \ \ \forall\ (a_j,b_j) \in T_j \qquad \text{if and only if} \qquad g_j(x) \le 0,$$

where

$$g_j(x) = \sup_{(a_j,b_j)\in T_j} \{a_j'x - b_j\}. \tag{1.27}$$

Thus, the robust linear programming problem (1.26) is equivalent to

$$\begin{aligned} &\text{minimize} \ \ c'x \\ &\text{subject to} \ \ g_j(x) \le 0, \qquad j=1,\ldots,r. \end{aligned}$$

For special choices of the set T_j, the function g_j can be expressed in closed form, and in the case where T_j is an ellipsoid, it turns out that the constraint $g_j(x) \le 0$ can be expressed in terms of a second order cone. To see this, let

$$T_j = \Big\{ (\overline{a}_j + P_j u_j, \overline{b}_j + q_j' u_j) \mid \|u_j\| \le 1,\, u_j \in \Re^{n_j} \Big\}, \tag{1.28}$$

where P_j is a given $n \times n_j$ matrix, $\overline{a}_j \in \Re^n$ and $q_j \in \Re^{n_j}$ are given vectors, and $\overline{b}_j$ is a given scalar. Then, from Eqs. (1.27) and (1.28),

$$\begin{aligned} g_j(x) &= \sup_{\|u_j\|\le 1} \Big\{ (\overline{a}_j + P_j u_j)'x - (\overline{b}_j + q_j' u_j) \Big\} \\ &= \sup_{\|u_j\|\le 1} (P_j'x - q_j)'u_j + \overline{a}_j'x - \overline{b}_j, \end{aligned}$$

and finally

$$g_j(x) = \|P_j'x - q_j\| + \overline{a}_j'x - \overline{b}_j.$$

Thus,

$$g_j(x) \le 0 \qquad \text{if and only if} \qquad (P_j'x - q_j, \overline{b}_j - \overline{a}_j'x) \in C_j,$$

where C_j is the second order cone of $\Re^{n_j+1}$; i.e., the "robust" constraint $g_j(x) \le 0$ is equivalent to a second order cone constraint. It follows that in the case of ellipsoidal uncertainty, the robust linear programming problem (1.26) is a SOCP of the form (1.24).

Example 1.2.4: (Quadratically Constrained Quadratic Problems)

Consider the quadratically constrained quadratic problem

$$\begin{aligned} &\text{minimize} \quad x'Q_0x + 2q_0'x + p_0 \\ &\text{subject to} \quad x'Q_jx + 2q_j'x + p_j \le 0, \quad j = 1, \ldots, r, \end{aligned}$$

where $Q_0, \ldots, Q_r$ are symmetric $n \times n$ positive definite matrices, $q_0, \ldots, q_r$ are vectors in $\Re^n$, and $p_0, \ldots, p_r$ are scalars. We show that the problem can be converted to the second order cone format. A similar conversion is also possible for the quadratic programming problem where Q_0 is positive definite and $Q_j = 0$, $j = 1, \ldots, r$.

Indeed, since each Q_j is symmetric and positive definite, we have

$$\begin{aligned} x'Q_jx + 2q_j'x + p_j &= \left(Q_j^{1/2}x\right)' Q_j^{1/2}x + 2\left(Q_j^{-1/2}q_j\right)' Q_j^{1/2}x + p_j \\ &= \|Q_j^{1/2}x + Q_j^{-1/2}q_j\|^2 + p_j - q_j'Q_j^{-1}q_j, \end{aligned}$$

for $j = 0, 1, \ldots, r$. Thus, the problem can be written as

$$\begin{aligned} &\text{minimize} \quad \|Q_0^{1/2}x + Q_0^{-1/2}q_0\|^2 + p_0 - q_0'Q_0^{-1}q_0 \\ &\text{subject to} \quad \|Q_j^{1/2}x + Q_j^{-1/2}q_j\|^2 + p_j - q_j'Q_j^{-1}q_j \le 0, \quad j = 1, \ldots, r, \end{aligned}$$

or, by neglecting the constant $p_0 - q_0'Q_0^{-1}q_0$,

$$\begin{aligned} &\text{minimize} \quad \|Q_0^{1/2}x + Q_0^{-1/2}q_0\| \\ &\text{subject to} \quad \|Q_j^{1/2}x + Q_j^{-1/2}q_j\| \le \left(q_j'Q_j^{-1}q_j - p_j\right)^{1/2}, \quad j = 1, \ldots, r. \end{aligned}$$

By introducing an auxiliary variable x_{n+1}, the problem can be written as

$$\begin{aligned} &\text{minimize} \quad x_{n+1} \\ &\text{subject to} \quad \|Q_0^{1/2}x + Q_0^{-1/2}q_0\| \le x_{n+1} \\ &\qquad\qquad \|Q_j^{1/2}x + Q_j^{-1/2}q_j\| \le \left(q_j'Q_j^{-1}q_j - p_j\right)^{1/2}, \quad j = 1, \ldots, r. \end{aligned}$$

It can be seen that this problem has the second order cone form (1.24). In particular, the first constraint is of the form $A_0x - b_0 \in C$, where C is the second order cone of $\Re^{n+1}$ and the $(n+1)$st component of $A_0x - b_0$ is x_{n+1}. The remaining r constraints are of the form $A_jx - b_j \in C$, where the $(n+1)$st component of $A_jx - b_j$ is the scalar $\left(q_j' Q_j^{-1} q_j - p_j\right)^{1/2}$.

We finally note that the problem of this example is special in that it has no duality gap, assuming its optimal value is finite, i.e., there is no need for the interior point conditions of Prop. 1.2.4. This can be traced to the fact that linear transformations preserve the closure of sets defined by quadratic constraints (see e.g., BNO03], Section 1.5.2).

1.2.3 Semidefinite Programming

In this section we consider the linear-conic problem (1.21) with C being the cone of matrices that are positive semidefinite.† This is called the *positive semidefinite cone*. To define the problem, we view the space of symmetric $n \times n$ matrices as the space $\Re^{n^2}$ with the inner product

$$< X, Y >= \text{trace}(XY) = \sum_{i=1}^{n} \sum_{j=1}^{n} x_{ij} y_{ij}.$$

The interior of C is the set of positive definite matrices.

The dual cone is

$$\hat{C} = \big\{Y \mid \text{trace}(XY) \geq 0,\ \forall\ X \in C\big\},$$

and it can be shown that $\hat{C} = C$, i.e., C is self-dual. Indeed, if $Y \notin C$, there exists a vector $v \in \Re^n$ such that

$$0 > v'Yv = \text{trace}(vv'Y).$$

Hence the positive semidefinite matrix $X = vv'$ satisfies $0 > \text{trace}(XY)$, so $Y \notin \hat{C}$ and it follows that $C \supset \hat{C}$. Conversely, let $Y \in C$, and let X be any positive semidefinite matrix. We can express X as

$$X = \sum_{i=1}^{n} \lambda_i e_i e_i',$$

where λ_i are the nonnegative eigenvalues of X, and e_i are corresponding orthonormal eigenvectors. Then,

$$\text{trace}(XY) = \text{trace}\left(Y \sum_{i=1}^{n} \lambda_i e_i e_i'\right) = \sum_{i=1}^{n} \lambda_i e_i' Y e_i \geq 0.$$

† As noted in Appendix A, throughout this book a positive semidefinite matrix is implicitly assumed to be symmetric.

It follows that $Y \in \hat{C}$ and $C \subset \hat{C}$. Thus C is self-dual, $C = \hat{C}$.

The semidefinite programming problem (SDP for short) is to minimize a linear function of a symmetric matrix over the intersection of an affine set with the positive semidefinite cone. It has the form

$$\begin{aligned} &\text{minimize} \quad < D, X > \\ &\text{subject to} \quad < A_i, X >= b_i, \quad i = 1, \ldots, m, \quad X \in C, \end{aligned} \tag{1.29}$$

where D, $A_1, \ldots, A_m$, are given $n \times n$ symmetric matrices, and $b_1, \ldots, b_m$, are given scalars. It is seen to be a special case of the primal problem in the left-hand side of the duality relation (1.21).

We can view the SDP as a problem with linear cost, linear constraints, and a convex set constraint. Then, similar to the case of SOCP, it can be verified that the dual problem (1.20), as given by the right-hand side of the duality relation (1.21), takes the form

$$\begin{aligned} &\text{maximize} \quad b'\lambda \\ &\text{subject to} \quad D - (\lambda_1 A_1 + \cdots + \lambda_m A_m) \in C, \end{aligned} \tag{1.30}$$

where $b = (b_1, \ldots, b_m)$ and the maximization is over the vector $\lambda = (\lambda_1, \ldots, \lambda_m)$. By applying the Linear-Conic Duality Theorem (Prop. 1.2.3), we have the following proposition.

Proposition 1.2.5: (Semidefinite Duality Theorem) Consider the primal SDP (1.29), and its dual problem (1.30).

(a) If the optimal value of the primal problem is finite and there exists a primal-feasible solution, which is positive definite, then there is no duality gap, and the dual problem has an optimal solution.

(b) If the optimal value of the dual problem is finite and there exist scalars $\overline{\lambda}_1, \ldots, \overline{\lambda}_m$ such that $D - (\overline{\lambda}_1 A_1 + \cdots + \overline{\lambda}_m A_m)$ is positive definite, then there is no duality gap, and the primal problem has an optimal solution.

The SDP is a fairly general problem. In particular, it can be shown that a SOCP can be cast as a SDP. Thus SDP involves a more general structure than SOCP. This is consistent with the practical observation that the latter problem is generally more amenable to computational solution. We provide some examples of problem formulation as an SDP.

Example 1.2.5: (Minimizing the Maximum Eigenvalue)

Given a symmetric $n \times n$ matrix $M(\lambda)$, which depends on a parameter vector $\lambda = (\lambda_1, \ldots, \lambda_m)$, we want to choose λ so as to minimize the maximum eigenvalue of $M(\lambda)$. We pose this problem as

$$\begin{array}{ll} \text{minimize} & z \\ \text{subject to} & \text{maximum eigenvalue of } M(\lambda) \leq z, \end{array}$$

or equivalently

$$\begin{array}{ll} \text{minimize} & z \\ \text{subject to} & zI - M(\lambda) \in C, \end{array}$$

where I is the $n \times n$ identity matrix, and C is the semidefinite cone. If $M(\lambda)$ is an affine function of λ,

$$M(\lambda) = M_0 + \lambda_1 M_1 + \cdots + \lambda_m M_m,$$

the problem has the form of the dual problem (1.30), with the optimization variables being $(z, \lambda_1, \ldots, \lambda_m)$.

Example 1.2.6: (Semidefinite Relaxation – Lower Bounds for Discrete Optimization Problems)

Semidefinite programming provides a means for deriving lower bounds to the optimal value of several types of discrete optimization problems. As an example, consider the following quadratic problem with quadratic equality constraints

$$\begin{array}{ll} \text{minimize} & x'Q_0 x + a_0' x + b_0 \\ \text{subject to} & x'Q_i x + a_i' x + b_i = 0, \quad i = 1, \ldots, m, \end{array} \tag{1.31}$$

where $Q_0, \ldots, Q_m$ are symmetric $n \times n$ matrices, $a_0, \ldots, a_m$ are vectors in $\Re^n$, and $b_0, \ldots, b_m$ are scalars.

This problem can be used to model broad classes of discrete optimization problems. To see this, consider an integer constraint that a variable x_i must be either 0 or 1. Such a constraint can be expressed by the quadratic equality $x_i^2 - x_i = 0$. Furthermore, a linear inequality constraint $a_j' x \leq b_j$ can be expressed as the quadratic equality constraint $y_j^2 + a_j' x - b_j = 0$, where y_j is an additional variable.

Introducing a multiplier vector $\lambda = (\lambda_1, \ldots, \lambda_m)$, the dual function is given by

$$q(\lambda) = \inf_{x \in \Re^n} \big\{ x'Q(\lambda)x + a(\lambda)'x + b(\lambda) \big\},$$

where

$$Q(\lambda) = Q_0 + \sum_{i=1}^m \lambda_i Q_i, \quad a(\lambda) = a_0 + \sum_{i=1}^m \lambda_i a_i, \quad b(\lambda) = b_0 + \sum_{i=1}^m \lambda_i b_i.$$

Let f^* and q^* be the optimal values of problem (1.31) and its dual, and note that by weak duality, we have $f^* \geq q^*$. By introducing an auxiliary scalar variable ξ, we see that the dual problem is to find a pair (ξ, λ) that solves the problem

$$\begin{aligned} &\text{maximize} \quad \xi \\ &\text{subject to} \quad q(\lambda) \geq \xi. \end{aligned}$$

The constraint $q(\lambda) \geq \xi$ of this problem can be written as

$$\inf_{x \in \Re^n} \big\{ x'Q(\lambda)x + a(\lambda)'x + b(\lambda) - \xi \big\} \geq 0,$$

or equivalently, introducing a scalar variable t and multiplying with t^2,

$$\inf_{x \in \Re^n,\, t \in \Re} \big\{ (tx)'Q(\lambda)(tx) + a(\lambda)'(tx)t + \big(b(\lambda) - \xi\big)t^2 \big\} \geq 0.$$

Writing $y = tx$, this relation takes the form of a quadratic in (y, t),

$$\inf_{y \in \Re^n,\, t \in \Re} \big\{ y'Q(\lambda)y + a(\lambda)'yt + \big(b(\lambda) - \xi\big)t^2 \big\} \geq 0,$$

or

$$\begin{pmatrix} Q(\lambda) & \frac{1}{2}a(\lambda) \\ \frac{1}{2}a(\lambda)' & b(\lambda) - \xi \end{pmatrix} \in C, \tag{1.32}$$

where C is the positive semidefinite cone. Thus the dual problem is equivalent to the SDP of maximizing ξ over all (ξ, λ) satisfying the constraint (1.32), and its optimal value q^* is a lower bound to f^*.

1.3 ADDITIVE COST PROBLEMS

In this section we focus on a structural characteristic that arises in several important contexts: a cost function f that is the sum of a large number of components $f_i : \Re^n \mapsto \Re$,

$$f(x) = \sum_{i=1}^{m} f_i(x). \tag{1.33}$$

Such cost functions can be minimized with specialized methods, called *incremental*, which exploit their additive structure, by updating x using one component function f_i at a time (see Section 2.1.5). Problems with additive cost functions can also be treated with specialized outer and inner linearization methods that approximate the component functions f_i individually (rather than approximating f); see Section 4.4.

An important special case is the cost function of the dual of a separable problem

$$\begin{aligned} &\text{maximize} \quad \sum_{i=1}^{m} q_i(\mu) \\ &\text{subject to} \quad \mu \geq 0, \end{aligned}$$

where

$$q_i(\mu) = \inf_{x_i \in X_i} \left\{ f_i(x_i) + \sum_{j=1}^{r} \mu_j g_{ij}(x_i) \right\},$$

and $\mu = (\mu_1, \ldots, \mu_r)$ [cf. Eq. (1.8)]. After a sign change to convert to minimization it takes the form (1.33) with $f_i(\mu) = -q_i(\mu)$. This is a major class of additive cost problems.

We will next describe some applications from a variety of fields. The following five examples arise in many machine learning contexts.

Example 1.3.1: (Regularized Regression)

This is a broad class of applications that relate to parameter estimation. The cost function involves a sum of terms $f_i(x)$, each corresponding to the error between some data and the output of a parametric model, with x being the vector of parameters. An example is linear least squares problems, also referred to as *linear regression* problems, where f_i has quadratic structure. Often a convex regularization function $R(x)$ is added to the least squares objective, to induce desirable properties of the solution and/or the corresponding algorithms. This gives rise to problems of the form

$$\begin{aligned} &\text{minimize} \quad R(x) + \tfrac{1}{2} \textstyle\sum_{i=1}^{m} (c_i'x - b_i)^2 \\ &\text{subject to} \ \ x \in \Re^n, \end{aligned}$$

where c_i and b_i are given vectors and scalars, respectively. The regularization function R is often taken to be differentiable, and particularly quadratic. However, there are practically important examples of nondifferentiable choices (see the next example).

In statistical applications, such a problem arises when constructing a linear model for an unknown input-output relation. The model involves a vector of parameters x, to be determined, which weigh input data (the components of the vectors c_i). The inner products $c_i'x$ produced by the model are matched against the scalars b_i, which are observed output data, corresponding to inputs c_i from the true input-output relation that we try to represent. The optimal vector of parameters x^* provides the model that (in the absence of a regularization function) minimizes the sum of the squared errors $(c_i'x^* - b_i)^2$.

In a more general version of the problem, a nonlinear parametric model is constructed, giving rise to a nonlinear least squares problem of the form

$$\begin{aligned} &\text{minimize} \quad R(x) + \sum_{i=1}^{m} \big|g_i(x)\big|^2 \\ &\text{subject to} \ \ x \in \Re^n, \end{aligned}$$

where $g_i : \Re^n \mapsto \Re$ are given nonlinear functions that depend on the data. This is also a common problem, referred to as *nonlinear regression*, which, however, is often nonconvex [it is convex if the functions g_i are convex *and* also nonnegative, i.e., $g_i(x) \geq 0$ for all $x \in \Re^n$].

It is also possible to use a nonquadratic function of the error between some data and the output of a linear parametric model. Thus in place of the squared error $(1/2)(c_i'x - b_i)^2$, we may use $h_i(c_i'x - b_i)$, where $h_i : \Re \mapsto \Re$ is a convex function, leading to the problem

$$\begin{aligned} &\text{minimize} \quad R(x) + \sum_{i=1}^{m} h_i(c_i'x - b_i) \\ &\text{subject to} \quad x \in \Re^n. \end{aligned}$$

Generally the choice of the function h_i is dictated by statistical modeling considerations, for which the reader may consult the relevant literature. An example is

$$h_i(c_i'x - b_i) = |c_i'x - b_i|,$$

which tends to result in a more robust estimate than least squares in the presence of large outliers in the data. This is known as the *least absolute deviations* method.

There are also constrained variants of the problems just discussed, where the parameter vector x is required to belong to some subset of $\Re^n$, such as the nonnegative orthant or a "box" formed by given upper and lower bounds on the components of x. Such constraints may be used to encode into the model some prior knowledge about the nature of the solution.

Example 1.3.2: (ℓ_1-Regularization)

A popular approach to regularized regression involves ℓ_1*-regularization*, where

$$R(x) = \gamma \|x\|_1 = \gamma \sum_{j=1}^{n} |x^j|,$$

γ is a positive scalar and x^j is the jth coordinate of x. The reason for the popularity of the ℓ_1 norm $\|x\|_1$ is that it tends to produce optimal solutions where a greater number of components x^j are zero, relative to the case of quadratic regularization (see Fig. 1.3.1). This is considered desirable in many statistical applications, where the number of parameters to include in a model may not be known a priori; see e.g., [Tib96], [DoE03], [BJM12]. The special case where a linear least squares model is used,

$$\begin{aligned} &\text{minimize} \quad \gamma \|x\|_1 + \tfrac{1}{2} \textstyle\sum_{i=1}^{m} (c_i'x - b_i)^2 \\ &\text{subject to} \quad x \in \Re^n, \end{aligned}$$

is known as the *lasso problem*.

In a generalization of the lasso problem, the ℓ_1 regularization function $\|x\|_1$ is replaced by a scaled version $\|Sx\|_1$, where S is some scaling matrix.

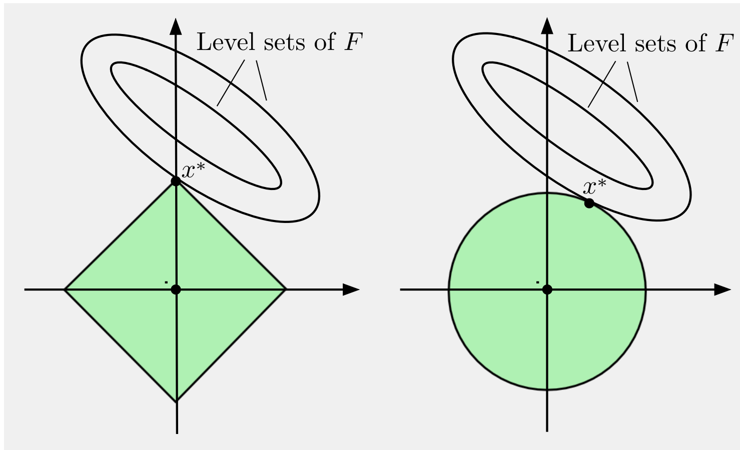

Figure 1.3.1. Illustration of the effect of ℓ_1-regularization for cost functions of the form $\gamma\|x\|_1+F(x)$, where $\gamma > 0$ and $F : \Re^n \mapsto \Re$ is differentiable (figure in the left-hand side). The optimal solution x^* tends to have more zero components than in the corresponding quadratic regularization case, illustrated in the right-hand side.

The term $\|Sx\|_1$ then induces a penalty on some undesirable characteristic of the solution. For example the problem

$$\begin{aligned} &\text{minimize} \quad \gamma \sum_{i=1}^{n-1} |x_{i+1} - x_i| + \tfrac{1}{2} \textstyle\sum_{i=1}^m (c_i'x - b_i)^2 \\ &\text{subject to} \quad x \in \Re^n, \end{aligned}$$

is known as the *total variation denoising problem*; see e.g., [ROF92], [Cha04], [BeT09a]. The regularization term here encourages consecutive variables to take similar values, and tends to produce more smoothly varying solutions.

Another related example is *matrix completion with nuclear norm regularization*; see e.g., [CaR09], [CaT10], [RFP10], [Rec11], [ReR13]. Here the minimization is over all $m \times n$ matrices X, with components denoted X_{ij}. We have a set of entries M_{ij}, $(i,j) \in \Omega$, where Ω is a subset of index pairs, and we want to find X whose entries X_{ij} are close to M_{ij} for $(i,j) \in \Omega$, and has as small rank as possible, a property that is desirable on the basis of statistical considerations. The following more tractable version of the problem is solved instead:

$$\begin{aligned} &\text{minimize} \quad \gamma\|X\|_* + \tfrac{1}{2} \textstyle\sum_{(i,j)\in\Omega} (X_{ij} - M_{ij})^2 \\ &\text{subject to} \quad X \in \Re^{m\times n}, \end{aligned}$$

where $\|X\|_*$ is the *nuclear norm* of X, defined as the sum of the singular values of X. There is substantial theory that justifies this approximation, for which we refer to the literature. It turns out that the nuclear norm is a convex function with some nice properties. In particular, its subdifferential at any X can be conveniently characterized for use in algorithms.

Let us finally note that sometimes additional regularization functions are used in conjunction with ℓ_1-type terms. An example is the sum of a quadratic and an ℓ_1-type term.

Example 1.3.3: (Classification)

In the regression problems of the preceding examples we aim to construct a parametric model that matches well an input-output relationship based on given data. Similar problems arise in a classification context, where we try to construct a parametric model for predicting whether an object with certain characteristics (also called features) belongs to a given category or not.

We assume that each object is characterized by a *feature vector* c that belongs to $\Re^n$ and a *label* b that takes the values $+1$ or -1, if the object belongs to the category or not, respectively. As illustration consider a credit card company that wishes to classify applicants as "low risk" (+1) or "high risk" (-1), with each customer characterized by n scalar features of financial and personal type.

We are given data, which is a set of feature-label pairs (c_i, b_i), $i = 1, \ldots, m$. Based on this data, we want to find a parameter vector $x \in \Re^n$ and a scalar $y \in \Re$ such that the sign of $c'x + y$ is a good predictor of the label of an object with feature vector c. Thus, loosely speaking, x and y should be such that for "most" of the given feature-label data (c_i, b_i) we have

$$c_i'x + y > 0, \qquad \text{if } b_i = +1,$$

$$c_i'x + y < 0, \qquad \text{if } b_i = -1.$$

In the statistical literature, $c'x + y$ is often called the *discriminant function*, and the value of

$$b_i(c_i'x + y),$$

for a given object i provides a measure of "margin" to misclassification of the object. In particular, a classification error is made for object i when $b_i(c_i'x + y) < 0$.

Thus it makes sense to formulate classification as an optimization problem where negative values of $b_i(c_i'x + y)$ are penalized. This leads to the problem

$$\begin{aligned} &\text{minimize} \quad R(x) + \sum_{i=1}^{m} h\big(b_i(c_i'x + y)\big) \\ &\text{subject to} \quad x \in \Re^n,\ y \in \Re, \end{aligned}$$

where R is a suitable regularization function, and $h : \Re \mapsto \Re$ is a convex function that penalizes negative values of its argument. It would make some sense to use a penalty of one unit for misclassification, i.e.,

$$h(z) = \begin{cases} 0 & \text{if } z \geq 0, \\ 1 & \text{if } z < 0, \end{cases}$$

but such a penalty function is discontinuous. To obtain a continuous cost function, we allow a continuous transition of h from negative to positive values, leading to a variety of nonincreasing functions h. The choice of h depends on the given application and other theoretical considerations for which we refer to the literature. Some common examples are

$$h(z) = e^{-z}, \qquad \text{(exponential loss)},$$

$$h(z) = \log\left(1 + e^{-z}\right), \qquad \text{(logistic loss)},$$

$$h(z) = \max\left\{0, 1 - z\right\}, \qquad \text{(hinge loss)}.$$

For the case of logistic loss the method comes under the methodology of *logistic regression*, and for the case of hinge loss the method comes under the methodology of *support vector machines*. As in the case of regression, the regularization function R could be quadratic, the ℓ_1 norm, or some scaled version or combination thereof. There is extensive literature on these methodologies and their applications, to which we refer for further discussion.

Example 1.3.4: (Nonnegative Matrix Factorization)

The nonnegative matrix factorization problem is to approximately factor a given nonnegative matrix B as CX, where C and X are nonnegative matrices to be determined via the optimization

$$\begin{array}{ll} \text{minimize} & \|CX - B\|_F^2 \\ \text{subject to} & C \geq 0, \; X \geq 0. \end{array}$$

Here $\|\cdot\|_F$ denotes the Frobenius norm of a matrix ($\|M\|_F^2$ is the sum of the squares of the scalar components of M). The matrices B, C, and X must have compatible dimensions, with the column dimension of C usually being much smaller than its row dimension, so that CX is a low-rank approximation of B. In some versions of the problem some of the nonnegativity constraints on the components of C and X may be relaxed. Moreover, regularization terms may be added to the cost function to induce sparsity or some other effect, similar to earlier examples in this section.

This problem, formulated in the 90s, [PaT94], [Paa97], [LeS99], has become a popular model for regression-type applications such as the ones of Example 1.3.1, but with the vectors c_i in the least squares objective $\sum_{i=1}^m (c_i'x - b_i)^2$ being unknown and subject to optimization. In the regression context of Example 1.3.1, we aim to (approximately) represent the data in the range space of the matrix C whose rows are the vectors c_i', and we may view C as a matrix of known basis functions. In the matrix factorization context of the present example, we aim to discover a "good" matrix C of basis functions that represents well the given data, i.e., the matrix B.

An important characteristic of the problem is that its cost function is not convex jointly in (C, X). However, it is convex in each of the matrices C and X individually, when the other matrix is held fixed. This facilitates the application of algorithms that involve alternate minimizations with respect to C and with respect to X; see Section 6.5. We refer to the literature, e.g., the papers [BBL07], [Lin07], [GoZ12], for a discussion of related algorithmic issues.

Example 1.3.5: (Maximum Likelihood Estimation)

The maximum likelihood approach is a major statistical inference methodology for parameter estimation, which is described in many sources (see e.g., the textbooks [Was04], [HTF09]). In fact in many cases, a maximum likelihood formulation is used to provide a probabilistic justification of the regression and classification models of the preceding examples.

Here we observe a sample of a random vector Z whose distribution $P_Z(\cdot;x)$ depends on an unknown parameter vector $x \in \Re^n$. For simplicity we assume that Z can take only a finite set of values, so that $P_Z(z;x)$ is the probability that Z takes the value z when the parameter vector has the value x. We estimate x based on the given sample value z, by solving the problem

$$\begin{aligned} &\text{maximize} \quad P_Z(z;x) \\ &\text{subject to} \;\; x \in \Re^n. \end{aligned} \tag{1.34}$$

The cost function $P_Z(z;\cdot)$ of this problem may either have an additive structure or may be equivalent to a problem that has an additive structure. For example the event that $Z = z$ may be the union of a large number of disjoint events, so $P_Z(z;x)$ is the sum of the probabilities of these events. For another important context, suppose that the data z consists of m independent samples $z_1,\ldots,z_m$ drawn from a distribution $P(\cdot;x)$, in which case

$$P_Z(z;x) = P(z_1;x)\cdots P(z_m;x).$$

Then the maximization (1.34) is equivalent to the additive cost minimization

$$\begin{aligned} &\text{minimize} \quad \sum_{i=1}^m f_i(x) \\ &\text{subject to} \;\; x \in \Re^n, \end{aligned}$$

where

$$f_i(x) = -\log P(z_i;x).$$

In many applications the number of samples m is very large, in which case special methods that exploit the additive structure of the cost are recommended. Often a suitable regularization term is added to the cost function, similar to the preceding examples.

Example 1.3.6: (Minimization of an Expected Value - Stochastic Programming)

An important context where additive cost functions arise is the minimization of an expected value

$$\begin{aligned} &\text{minimize} \quad E\big\{F(x,w)\big\} \\ &\text{subject to} \;\; x \in X, \end{aligned}$$

where w is a random variable taking a finite but very large number of values w_i, $i = 1, \ldots, m$, with corresponding probabilities π_i. Then the cost function consists of the sum of the m functions $\pi_i F(x, w_i)$.

For example, in *stochastic programming*, a classical model of two-stage optimization under uncertainty, a vector $x \in X$ is selected, a random event occurs that has m possible outcomes $w_1, \ldots, w_m$, and another vector $y \in Y$ is selected with knowledge of the outcome that occurred (see e.g., the books [BiL97], [KaW94], [Pre95], [SDR09]). Then for optimization purposes, we need to specify a different vector $y_i \in Y$ for each outcome w_i. The problem is to minimize the expected cost

$$F(x) + \sum_{i=1}^{m} \pi_i G_i(y_i),$$

where $G_i(y_i)$ is the cost associated with the choice y_i and the occurrence of w_i, and π_i is the corresponding probability. This is a problem with an additive cost function.

Additive cost functions also arise when the expected value cost function $E\big\{F(x, w)\big\}$ is approximated by an m-sample average

$$f(x) = \frac{1}{m} \sum_{i=1}^{m} F(x, w_i),$$

where w_i are independent samples of the random variable w. The minimum of the sample average $f(x)$ is then taken as an approximation of the minimum of $E\big\{F(x, w)\big\}$.

Generally additive cost problems arise when we want to strike a balance between several types of costs by lumping them into a single cost function. The following is an example of a different character than the preceding ones.

Example 1.3.7: (Weber Problem in Location Theory)

A basic problem in location theory is to find a point x in the plane whose sum of weighted distances from a given set of points $y_1, \ldots, y_m$ is minimized. Mathematically, the problem is

$$\begin{aligned} &\text{minimize} \quad \sum_{i=1}^{m} w_i \|x - y_i\| \\ &\text{subject to} \quad x \in \Re^n, \end{aligned}$$

where $w_1, \ldots, w_m$ are given positive scalars. This problem has many variations, including constrained versions, and descends from the famous Fermat-Torricelli-Viviani problem (see [BMS99] for an account of the history of this problem). We refer to the book [DrH04] for a survey of recent research, and to the paper [BeT10] for a discussion that is relevant to our context.

The structure of the additive cost function (1.33) often facilitates the use of a distributed computing system that is well-suited for the incremental approach. The following is an illustrative example.

Example 1.3.8: (Distributed Incremental Optimization – Sensor Networks)

Consider a network of m sensors where data are collected and are used to solve some inference problem involving a parameter vector x. If $f_i(x)$ represents an error penalty for the data collected by the ith sensor, the inference problem involves an additive cost function $\sum_{i=1}^m f_i$. While it is possible to collect all the data at a fusion center where the problem will be solved in centralized manner, it may be preferable to adopt a distributed approach in order to save in data communication overhead and/or take advantage of parallelism in computation. In such an approach the current iterate x_k is passed on from one sensor to another, with each sensor i performing an incremental iteration involving just its local component f_i. The entire cost function need not be known at any one location. For further discussion we refer to representative sources such as [RaN04], [RaN05], [BHG08], [MRS10], [GSW12], and [Say14].

The approach of computing incrementally the values and subgradients of the components f_i in a distributed manner can be substantially extended to apply to general systems of asynchronous distributed computation, where the components are processed at the nodes of a computing network, and the results are suitably combined [NBB01] (see our discussion in Sections 2.1.5 and 2.1.6).

Let us finally note a constrained version of additive cost problems where the functions f_i are extended real-valued. This is essentially equivalent to constraining x to lie in the intersection of the domains

$$X_i = \text{dom}(f_i),$$

resulting in a problem of the form

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^m f_i(x) \\ &\text{subject to} && x \in \cap_{i=1}^m X_i, \end{aligned}$$

where each f_i is real-valued over the set X_i. Methods that are well-suited for the unconstrained version of the problem where $X_i \equiv \Re^n$ can often be modified to apply to the constrained version, as we will see in Chapter 6, where we will discuss incremental constraint projection methods. However, the case of constraint sets with many components arises independently of whether the cost function is additive or not, and has its own character, as we discuss in the next section.

1.4 LARGE NUMBER OF CONSTRAINTS

In this section we consider problems of the form

$$\begin{aligned} &\text{minimize} \quad f(x) \\ &\text{subject to} \quad x \in X, \quad g_j(x) \leq 0, \quad j = 1, \ldots, r, \end{aligned} \tag{1.35}$$

where the number r of constraints is very large. Problems of this type occur often in practice, either directly or via reformulation from other problems. A similar type of problem arises when the abstract constraint set X consists of the intersection of many simpler sets:

$$X = \cap_{\ell \in L} X_\ell,$$

where L is a finite or infinite index set. There may or may not be additional inequality constraints $g_j(x) \leq 0$ like the ones in problem (1.35). We provide a few examples.

Example 1.4.1: (Feasibility and Minimum Distance Problems)

A simple but important problem, which arises in many contexts and embodies important algorithmic ideas, is a classical *feasibility problem*, where the objective is to find a common point within a collection of sets X_ℓ, $\ell \in L$, where each X_ℓ is a closed convex set. In the feasibility problem the cost function is zero. A somewhat more complex problem with a similar structure arises when there is a cost function, i.e., a problem of the form

$$\begin{aligned} &\text{minimize} \quad f(x) \\ &\text{subject to} \quad x \in \cap_{\ell \in L} X_\ell, \end{aligned}$$

where $f : \Re^n \mapsto \Re$. An important example is the minimum distance problem, where

$$f(x) = \|x - z\|,$$

for a given vector z and some norm $\|\cdot\|$. The following example is a special case.

Example 1.4.2: (Basis Pursuit)

Consider the problem

$$\begin{aligned} &\text{minimize} \quad \|x\|_1 \\ &\text{subject to} \quad Ax = b, \end{aligned} \tag{1.36}$$

where $\|\cdot\|_1$ is the ℓ_1 norm in $\Re^n$, A is a given $m \times n$ matrix, and b is a vector in $\Re^m$ that consists of m given measurements. We are trying to construct a linear model of the form $Ax = b$, where x is a vector of n scalar weights for a large number n of basis functions ($m < n$). We want to satisfy exactly the measurement equations $Ax = b$, while using only a few of the basis functions in our model. Consequently, we introduce the ℓ_1 norm in the cost function of problem (1.36), aiming to delineate a small subset of basis functions, corresponding to nonzero coordinates of x at the optimal solution. This is called the *basis pursuit* problem (see, e.g., [CDS01], [VaF08]), and its underlying idea is similar to the one of ℓ_1-regularization (cf. Example 1.3.2).

It is also possible to consider a norm other than ℓ_1 in Eq. (1.36). An example is the *atomic norm* $\|\cdot\|_{\mathcal{A}}$ induced by a subset $\mathcal{A}$ that is centrally symmetric around the origin ($a \in \mathcal{A}$ if and only if $-a \in \mathcal{A}$):

$$\|x\|_{\mathcal{A}} = \inf\big\{t > 0 \mid x \in t \cdot \mathrm{conv}(\mathcal{A})\big\}.$$

This problem, and other related problems involving atomic norms, have many applications; see for example [CRP12], [SBT12], [RSW13].

A related problem is

$$\begin{aligned} &\text{minimize} \quad \|X\|_* \\ &\text{subject to} \quad AX = B, \end{aligned}$$

where the optimization is over all $m \times n$ matrices X. The matrices A, B are given and have dimensions $\ell \times m$ and $\ell \times n$, respectively, and $\|X\|_*$ is the nuclear norm of X. This problem aims to produce a low-rank matrix X that satisfies an underdetermined set of linear equations $AX = B$ (see e.g., [CaR09], [RFP10], [RXB11]). When these equations specify that a subset of entries X_{ij}, $(i,j) \in \Omega$, are fixed at given values M_{ij},

$$X_{ij} = M_{ij}, \qquad (i,j) \in \Omega,$$

we obtain an alternative formulation of the matrix completion problem discussed in Example 1.3.2.

Example 1.4.3: (Minimax Problems)

In a minimax problem the cost function has the form

$$f(x) = \sup_{z \in Z} \phi(x, z),$$

where Z is a subset of some space and $\phi(\cdot, z)$ is a real-valued function for each $z \in Z$. We want to minimize f subject to $x \in X$, where X is a given constraint set. By introducing an artificial scalar variable y, we may transform such a problem to the general form

$$\begin{aligned} &\text{minimize} \quad y \\ &\text{subject to} \quad x \in X, \quad \phi(x, z) \le y, \ \ \forall\, z \in Z, \end{aligned}$$

which involves a large number of constraints (one constraint for each z in the set Z, which could be infinite). Of course in this problem the set X may also be of the form $X = \cap_{\ell \in L} X_\ell$ as in earlier examples.

Example 1.4.4: (Basis Function Approximation for Separable Problems – Approximate Dynamic Programming)

Let us consider a large-scale separable problem of the form

$$\begin{aligned} &\text{minimize} \quad \sum_{i=1}^{m} f_i(y_i) \\ &\text{subject to} \quad \sum_{i=1}^{m} g_{ij}(y_i) \leq 0, \quad \forall\, j = 1, \dots, r, \quad y \geq 0, \end{aligned} \tag{1.37}$$

where $f_i : \Re \mapsto \Re$ are scalar functions, and the dimension m of the vector $y = (y_1, \dots, y_m)$ is very large. One possible way to address this problem is to approximate y with a vector of the form Φx, where Φ is an $m \times n$ matrix. The columns of Φ may be relatively few, and may be viewed as basis functions for a low-dimensional approximation subspace $\{\Phi x \mid x \in \Re^n\}$. We replace problem (1.37) with the approximate version

$$\begin{aligned} &\text{minimize} \quad \sum_{i=1}^{m} f_i(\phi_i' x) \\ &\text{subject to} \quad \sum_{i=1}^{m} g_{ij}(\phi_i' x) \leq 0, \quad \forall\, j = 1, \dots, r, \\ &\qquad\qquad\quad \phi_i' x \geq 0, \quad i = 1, \dots, m, \end{aligned} \tag{1.38}$$

where ϕ_i' denotes the ith row of Φ, and $\phi_i' x$ is viewed as an approximation of y_i. Thus the dimension of the problem is reduced from m to n. However, the constraint set of the problem became more complicated, because the simple constraints $y_i \geq 0$ take the more complex form $\phi_i' x \geq 0$. Moreover the number m of additive components in the cost function, as well as the number of its constraints is still large. Thus the problem has the additive cost structure of the preceding section, as well as a large number of constraints.

An important application of this approach is in approximate dynamic programming (see e.g., [BeT96], [SuB98], [Pow11], [Ber12]), where the functions f_i and g_{ij} are linear. The corresponding problem (1.37) relates to the solution of the optimality condition (Bellman equation) of an infinite horizon Markovian decision problem (the constraint $y \geq 0$ may not be present in this context). Here the numbers m and r are often astronomical (in fact r can be much larger than m), in which case an exact solution cannot be obtained. For such problems, approximation based on problem (1.38) has been one of the major algorithmic approaches (see [Ber12] for a textbook presentation and references). For very large m, it may be impossible to calculate the cost function value $\sum_{i=1}^{m} f_i(\phi_i' x)$ for a given x, and one may at most be able to sample individual cost components f_i. For this reason optimization by stochastic simulation is one of the most prominent approaches in large scale dynamic programming.

Let us also mention that related approaches based on randomization and simulation have been proposed for the solution of large scale instances of classical linear algebra problems; see [BeY09], [Ber12] (Section 7.3), [DMM06], [StV09], [HMT10], [Nee10], [DMM11], [WaB13a], [WaB13b].

A large number of constraints also arises often in problems involving a graph, and may be handled with algorithms that take into account the graph structure. The following example is typical.

Example 1.4.5: (Optimal Routing in a Network – Multicommodity Flows)

Consider a directed graph that is used to transfer "commodities" from given supply points to given demand points. We are given a set W of ordered node pairs $w = (i, j)$. The nodes i and j are referred to as the *origin* and the *destination* of w, respectively, and w is referred to as an OD pair. For each w, we are given a scalar r_w referred to as the *input* of w. For example, in the context of routing of data in a communication network, r_w (measured in data units/second) is the arrival rate of traffic entering and exiting the network at the origin and the destination of w, respectively. The objective is to divide each r_w among the many paths from origin to destination in a way that the resulting total arc flow pattern minimizes a suitable cost function.

We denote:

P_w: A given set of paths that start at the origin and end at the destination of w. All arcs on each of these paths are oriented in the direction from the origin to the destination.

x_p: The portion of r_w assigned to path p, also called the *flow of path p*.

The collection of all path flows $\{x_p \mid p \in P_w, w \in W\}$ must satisfy the constraints

$$\sum_{p \in P_w} x_p = r_w, \qquad \forall\, w \in W, \tag{1.39}$$

$$x_p \geq 0, \qquad \forall\, p \in P_w,\ w \in W. \tag{1.40}$$

The total flow F_{ij} of arc (i, j) is the sum of all path flows traversing the arc:

$$F_{ij} = \sum_{\substack{\text{all paths } p \\ \text{containing } (i,j)}} x_p. \tag{1.41}$$

Consider a cost function of the form

$$\sum_{(i,j)} D_{ij}(F_{ij}). \tag{1.42}$$

The problem is to find a set of path flows $\{x_p\}$ that minimize this cost function subject to the constraints of Eqs. (1.39)-(1.41). It is typically assumed that D_{ij} is a convex function of F_{ij}. In data routing applications, the form of D_{ij} is often based on a queueing model of average delay, in which case D_{ij} is continuously differentiable within its domain (see e.g., [BeG92]). In a related context, arising in optical networks, the problem involves additional integer constraints on x_p, but may be addressed as a problem with continuous flow variables (see [OzB03]).

The preceding problem is known as a *multicommodity network flow problem*. The terminology reflects the fact that the arc flows consist of several different commodities; in the present example the different commodities are the data of the distinct OD pairs. This problem also arises in essentially identical form in traffic network equilibrium problems (see e.g., [FlH95], [Ber98], [Ber99], [Pat99], [Pat04]). The special case where all OD pairs have the same end node, or all OD pairs have the same start node, is known as the *single commodity network flow problem*, a much easier type of problem, for which there are efficient specialized algorithms that tend to be much faster than their multicommodity counterparts (see textbooks such as [Ber91], [Ber98]).

By expressing the total flows F_{ij} in terms of the path flows in the cost function (1.42) [using Eq. (1.41)], the problem can be formulated in terms of the path flow variables $\{x_p \mid p \in P_w,\ w \in W\}$ as

$$\begin{aligned} &\text{minimize} \ \ D(x) \\ &\text{subject to} \ \ \sum_{p \in P_w} x_p = r_w, \quad \forall\, w \in W, \\ &x_p \geq 0, \quad \forall\, p \in P_w,\ w \in W, \end{aligned}$$

where

$$D(x) = \sum_{(i,j)} D_{ij} \left(\sum_{\substack{\text{all paths } p \\ \text{containing } (i,j)}} x_p \right)$$

and x is the vector of path flows x_p. There is a potentially huge number of variables as well as constraints in this problem. However, by judiciously taking into account the special structure of the problem, the constraint set can be simplified and approximated by the convex hull of a small number of vectors x, and the number of variables and constraints can be reduced to a manageable size (see e.g., [BeG83], [FlH95], [OMV00], and our discussion in Section 4.2).

There are several approaches to handle a large number of constraints. One possibility, which points the way to some major classes of algorithms, is to initially discard some of the constraints, solve the corresponding less constrained problem, and later selectively reintroduce constraints that seem to be violated at the optimum. In Chapters 4-6, we will discuss methods of this type in some detail.

Another possibility is to replace constraints with penalties that assign high cost for their violation. In particular, we may replace problem (1.35) with

$$\begin{aligned} &\text{minimize} \quad f(x) + c \sum_{j=1}^{r} P\big(g_j(x)\big) \\ &\text{subject to} \ \ x \in X, \end{aligned}$$

where $P(\cdot)$ is a scalar penalty function satisfying $P(u) = 0$ if $u \leq 0$, and $P(u) > 0$ if $u > 0$, and c is a positive penalty parameter. We discuss this possibility in the next section.

1.5 EXACT PENALTY FUNCTIONS

In this section we discuss a transformation that is often useful in the context of constrained optimization algorithms. We will derive a form of equivalence between a constrained convex optimization problem, and a penalized problem that is less constrained or is entirely unconstrained. The motivation is that some convex optimization algorithms do not have constrained counterparts, but can be applied to a penalized unconstrained problem. Furthermore, in some analytical contexts, it is useful to be able to work with an equivalent problem that is less constrained.

We consider the convex programming problem

$$\begin{aligned} &\text{minimize} \quad f(x) \\ &\text{subject to} \quad x \in X, \quad g_j(x) \leq 0, \quad j = 1, \dots, r, \end{aligned} \tag{1.43}$$

where X is a convex subset of $\Re^n$, and $f : X \to \Re$ and $g_j : X \to \Re$ are given convex functions. We denote by f^* the primal optimal value, and by q^* the dual optimal value, i.e.,

$$q^* = \sup_{\mu \geq 0} q(\mu),$$

where

$$q(\mu) = \inf_{x \in X} \bigl\{ f(x) + \mu' g(x) \bigr\}, \qquad \forall\ \mu \geq 0,$$

with $g(x) = \bigl(g_1(x), \dots, g_r(x)\bigr)'$. We assume that $-\infty < q^* = f^* < \infty$.

We introduce a convex penalty function $P : \Re^r \mapsto \Re$, which satisfies

$$P(u) = 0, \qquad \forall\ u \leq 0, \tag{1.44}$$

$$P(u) > 0, \qquad \text{if } u_j > 0 \text{ for some } j = 1, \dots, r. \tag{1.45}$$

We consider solving in place of the original problem (1.43), the "penalized" problem

$$\begin{aligned} &\text{minimize} \quad f(x) + P\bigl(g(x)\bigr) \\ &\text{subject to} \quad x \in X, \end{aligned} \tag{1.46}$$

where the inequality constraints have been replaced by the extra cost $P\bigl(g(x)\bigr)$ for their violation. Some interesting examples of penalty functions are based on the squared or the absolute value of constraint violation:

$$P(u) = \frac{c}{2} \sum_{j=1}^{r} \bigl(\max\{0, u_j\}\bigr)^2,$$

and

$$P(u) = c \sum_{j=1}^{r} \max\{0, u_j\},$$

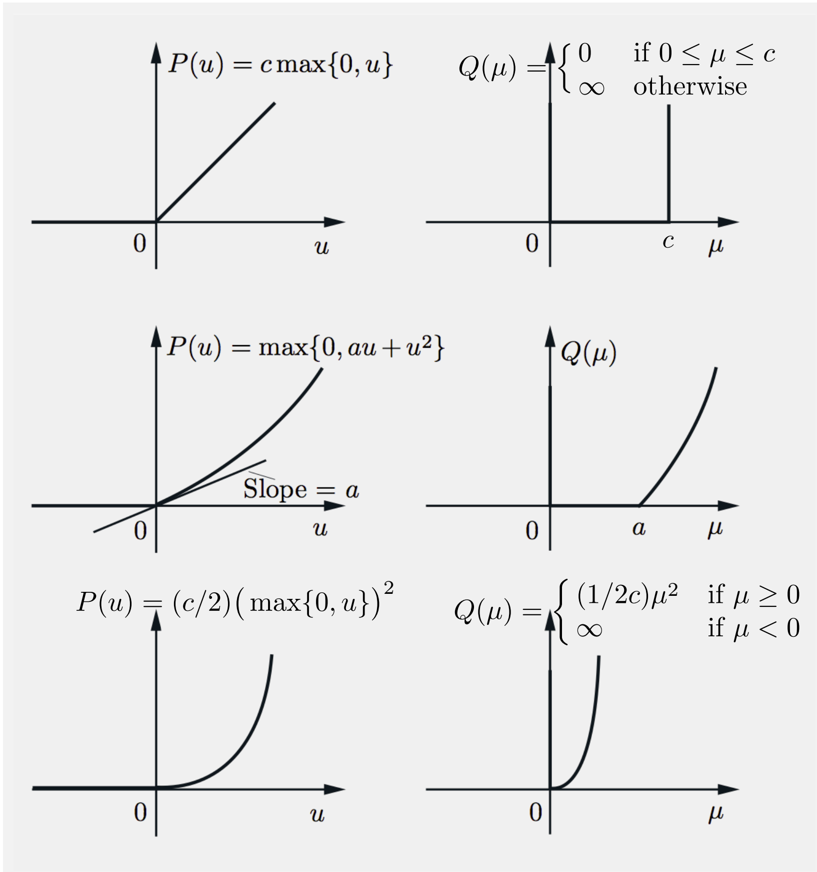

Figure 1.5.1. Illustration of various penalty functions P and their conjugate functions, denoted by Q. Because $P(u) = 0$ for $u \le 0$, we have $Q(\mu) = \infty$ for μ outside the nonnegative orthant.

where c is a positive penalty parameter. However, there are other possibilities that may be well-matched with the problem at hand.

The conjugate function of P is given by

$$Q(\mu) = \sup_{u \in \Re^r} \big\{ u'\mu - P(u) \big\},$$

and it can be seen that

$$Q(\mu) \ge 0, \qquad \forall\, \mu \in \Re^r,$$

$$Q(\mu) = \infty, \qquad \text{if } \mu_j < 0 \text{ for some } j = 1, \ldots, r.$$

Figure 1.5.1 shows some examples of one-dimensional penalty functions P, together with their conjugates.

Consider the primal function of the original constrained problem,

$$p(u) = \inf_{x \in X,\, g(x) \le u} f(x), \qquad u \in \Re^r.$$

We have,

$$\begin{aligned}
\inf_{x \in X} \big\{ f(x) + P\big(g(x)\big) \big\} &= \inf_{x \in X} \inf_{u \in \Re^r,\, g(x) \le u} \big\{ f(x) + P\big(g(x)\big) \big\} \\
&= \inf_{x \in X} \inf_{u \in \Re^r,\, g(x) \le u} \big\{ f(x) + P(u) \big\} \\
&= \inf_{x \in X,\, u \in \Re^r,\, g(x) \le u} \big\{ f(x) + P(u) \big\} \\
&= \inf_{u \in \Re^r} \inf_{x \in X,\, g(x) \le u} \big\{ f(x) + P(u) \big\} \\
&= \inf_{u \in \Re^r} \big\{ p(u) + P(u) \big\},
\end{aligned}$$

where for the second equality, we use the monotonicity relation†

$$u \le v \qquad \Rightarrow \qquad P(u) \le P(v).$$

Moreover, $-\infty < q^*$ and $f^* < \infty$ by assumption, and since for any μ with $q(\mu) > -\infty$, we have

$$p(u) \ge q(\mu) - \mu' u > -\infty, \qquad \forall\, u \in \Re^r,$$

it follows that $p(0) < \infty$ and $p(u) > -\infty$ for all $u \in \Re^r$, so p is proper.

We can now apply the Fenchel Duality Theorem (Prop. 1.2.1) with the identifications $f_1 = p$, $f_2 = P$, and $A = I$. We use the conjugacy relation between the primal function p and the dual function q to write

$$\inf_{u \in \Re^r} \big\{ p(u) + P(u) \big\} = \sup_{\mu \ge 0} \big\{ q(\mu) - Q(\mu) \big\}, \tag{1.47}$$

so that

$$\inf_{x \in X} \big\{ f(x) + P\big(g(x)\big) \big\} = \sup_{\mu \ge 0} \big\{ q(\mu) - Q(\mu) \big\}; \tag{1.48}$$

see Fig. 1.5.2. Note that the conditions for application of the theorem are satisfied since the penalty function P is real-valued, so that the relative

† To show this relation, we argue by contradiction. If there exist u and v with $u \le v$ and $P(u) > P(v)$, then by continuity of P, there must exist $\overline{u}$ close enough to u such that $\overline{u} < v$ and $P(\overline{u}) > P(v)$. Since P is convex, it is monotonically increasing along the halfline $\big\{ \overline{u} + \alpha(\overline{u} - v) \mid \alpha \ge 0 \big\}$, and since $P(\overline{u}) > P(v) \ge 0$, P takes positive values along this halfline. However, since $\overline{u} < v$, this halfline eventually enters the negative orthant, where P takes the value 0 by Eq. (1.44), a contradiction.

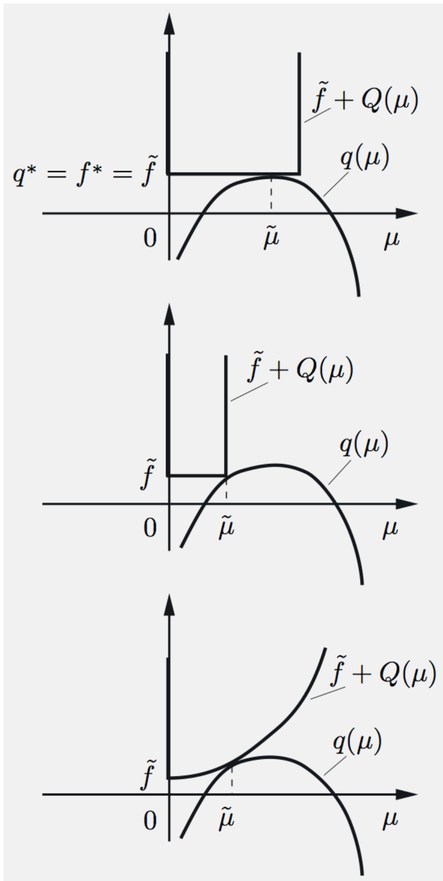

Figure 1.5.2. Illustration of the duality relation (1.48), and the optimal values of the penalized and the dual problem. Here f^* is the optimal value of the original problem, which is assumed to be equal to the optimal dual value q^*, while $\tilde{f}$ is the optimal value of the penalized problem,

$$\tilde{f} = \inf_{x \in X} \Big\{ f(x) + P\big(g(x)\big) \Big\}.$$

The point of contact of the graphs of the functions $\tilde{f} + Q(\mu)$ and $q(\mu)$ corresponds to the vector $\tilde{\mu}$ that attains the maximum in the relation

$$\tilde{f} = \max_{\mu \geq 0} \big\{ q(\mu) - Q(\mu) \big\}.$$

interiors of $\text{dom}(p)$ and $\text{dom}(P)$ have nonempty intersection. Furthermore, as part of the conclusions of part (a) of the Fenchel Duality Theorem, it follows that the supremum over $\mu \geq 0$ in Eq. (1.48) is attained.

Figure 1.5.2 suggests that in order for the penalized problem (1.46) to have the same optimal value as the original constrained problem (1.43), the conjugate Q must be “sufficiently flat” so that it is minimized by some dual optimal solution μ^*. This can be interpreted in terms of properties of subgradients, which are stated in Appendix B, Section 5.4: we must have $0 \in \partial Q(\mu^*)$ for some dual optimal solution μ^*, which by Prop. 5.4.3 in Appendix B, is equivalent to $\mu^* \in \partial P(0)$. This is part (a) of the following proposition, which was given in [Ber75a]. Parts (b) and (c) of the proposition deal with issues of equality of corresponding optimal solutions. The proposition assumes the convexity and other assumptions made in the early part in this section regarding problem (1.43) and the penalty function P.

Proposition 1.5.1: Consider problem (1.43), where we assume that $-\infty < q^* = f^* < \infty$.

(a) The penalized problem (1.46) and the original constrained problem (1.43) have equal optimal values if and only if there exists a dual optimal solution μ^* such that $\mu^* \in \partial P(0)$.

(b) In order for some optimal solution of the penalized problem (1.46) to be an optimal solution of the constrained problem (1.43), it is necessary that there exists a dual optimal solution μ^* such that

$$u'\mu^* \le P(u), \qquad \forall\ u \in \Re^r. \tag{1.49}$$

(c) In order for the penalized problem (1.46) and the constrained problem (1.43) to have the same set of optimal solutions, it is sufficient that there exists a dual optimal solution μ^* such that

$$u'\mu^* < P(u), \qquad \forall\ u \in \Re^r \text{ with } u_j > 0 \text{ for some } j. \tag{1.50}$$

Proof: (a) We have using Eqs. (1.47) and (1.48),

$$p(0) \ge \inf_{u\in\Re^r} \big\{p(u) + P(u)\big\} = \sup_{\mu\ge 0}\big\{q(\mu) - Q(\mu)\big\} = \inf_{x\in X} \Big\{f(x) + P\big(g(x)\big)\Big\}. \tag{1.51}$$

Since $f^* = p(0)$, we have

$$f^* = \inf_{x\in X} \Big\{f(x) + P\big(g(x)\big)\Big\}$$

if and only if equality holds in Eq. (1.51). This is true if and only if

$$0 \in \arg\min_{u\in\Re^r} \big\{p(u) + P(u)\big\},$$

which by Prop. 5.4.7 in Appendix B, is true if and only if there exists some $\mu^* \in -\partial p(0)$ with $\mu^* \in \partial P(0)$ (in view of the fact that P is real-valued). Since the set of dual optimal solutions is $-\partial p(0)$ (under our assumption $-\infty < q^* = f^* < \infty$; see Example 5.4.2, [Ber09]), the result follows.

(b) If x^* is an optimal solution of both problems (1.43) and (1.46), then by feasibility of x^*, we have $P\big(g(x^*)\big) = 0$, so these two problems have equal optimal values. From part (a), there must exist a dual optimal solution $\mu^* \in \partial P(0)$, which is equivalent to Eq. (1.49), by the subgradient inequality.

(c) If x^* is an optimal solution of the constrained problem (1.43), then $P\big(g(x^*)\big) = 0$, so we have

$$f^* = f(x^*) = f(x^*) + P\big(g(x^*)\big) \ge \inf_{x\in X}\Big\{f(x) + P\big(g(x)\big)\Big\}.$$

The condition (1.50) implies the condition (1.49), so that by part (a), equality holds throughout in the above relation, showing that x^* is also an optimal solution of the penalized problem (1.46).

Conversely, let $x^* \in X$ be an optimal solution of the penalized problem (1.46). If x^* is feasible [i.e., satisfies in addition $g(x^*) \leq 0$], then it is an optimal solution of the constrained problem (1.43) [since $P\big(g(x)\big) = 0$ for all feasible vectors x], and we are done. Otherwise x^* is infeasible in which case $g_j(x^*) > 0$ for some j. Then, by using the given condition (1.50), it follows that there exists a dual optimal solution μ^* and an $\epsilon > 0$ such that

$$\mu^{*\prime} g(x^*) + \epsilon < P\big(g(x^*)\big).$$

Let $\tilde{x}$ be a feasible vector such that $f(\tilde{x}) \leq f^* + \epsilon$. Since $P\big(g(\tilde{x})\big) = 0$ and $f^* = \min_{x \in X}\big\{f(x) + \mu^{*\prime} g(x)\big\}$, we obtain

$$f(\tilde{x}) + P\big(g(\tilde{x})\big) = f(\tilde{x}) \leq f^* + \epsilon \leq f(x^*) + \mu^{*\prime} g(x^*) + \epsilon.$$

By combining the last two relations, it follows that

$$f(\tilde{x}) + P\big(g(\tilde{x})\big) < f(x^*) + P\big(g(x^*)\big),$$

which contradicts the hypothesis that x^* is an optimal solution of the penalized problem (1.46). This completes the proof. **Q.E.D.**

As an illustration, consider the minimization of $f(x) = -x$ over all $x \in X = \{x \mid x \geq 0\}$ with $g(x) = x \leq 0$. The dual function is

$$q(\mu) = \inf_{x \geq 0} (\mu - 1)x, \qquad \mu \geq 0,$$

so $q(\mu) = 0$ for $\mu \in [1, \infty)$ and $q(\mu) = -\infty$ otherwise. Let $P(u) = c \max\{0, u\}$, so the penalized problem is $\min_{x \geq 0}\big\{-x + c\max\{0, x\}\big\}$. Then parts (a) and (b) of the proposition apply if $c \geq 1$. However, part (c) applies only if $c > 1$. In terms of Fig. 1.5.2, the conjugate of P is $Q(\mu) = 0$ if $\mu \in [0, c]$ and $Q(\mu) = \infty$ otherwise, so when $c = 1$, Q is "flat" over an area not including an interior point of the dual optimal solution set $[1, \infty)$.

To elaborate on the idea of the preceding example, let

$$P(u) = c \sum_{j=1}^{r} \max\{0, u_j\},$$

where $c > 0$. The condition $\mu^* \in \partial P(0)$, or equivalently,

$$u' \mu^* \leq P(u), \qquad \forall\, u \in \Re^r$$

[cf. Eq. (1.49)], is equivalent to

$$\mu_j^* \leq c, \qquad \forall\, j = 1, \ldots, r.$$

Similarly, the condition $u'\mu^* < P(u)$ for all $u \in \Re^r$ with $u_j > 0$ for some j [cf. Eq. (1.50)], is equivalent to

$$\mu_j^* < c, \qquad \forall\ j = 1, \ldots, r.$$

The reader may consult the literature for other results on exact penalty functions, starting with their first proposal in the book [Zan69]. The preceding development is based on [Ber75], and focuses on convex programming problems. For additional representative references, some of which also discuss nonconvex problems, see [HaM79], [Ber82a], [Bur91], [FeM91], [BNO03], [FrT07]. In what follows we develop an exact penalty function result for the case of an abstract constraint set, which will be used in the context of incremental constraint projection algorithms in Section 6.4.4.

A Distance-Based Exact Penalty Function

Let us discuss the case of a general Lipschitz continuous (not necessarily convex) cost function and an abstract constraint set $X \subset \Re^n$. The idea is to use a penalty that is proportional to the distance from X:

$$\text{dist}(x; X) = \inf_{y \in X} \|x - y\|.$$

The next proposition from [Ber11] provides the basic result (see Fig. 1.5.3).

Proposition 1.5.2: Let $f : \Re^n \mapsto \Re$ be a function that is Lipschitz continuous with constant L over a set $Y \subset \Re^n$, i.e.,

$$\big|f(x) - f(y)\big| \le L\|x - y\|, \qquad \forall\ x, y \in Y. \tag{1.52}$$

Let also X be a nonempty closed subset of Y, and let c be a scalar with $c > L$. Then x^* minimizes f over X if and only if x^* minimizes

$$F_c(x) = f(x) + c\,\text{dist}(x; X)$$

over Y.

Proof: For any $x \in Y$, let $\hat{x}$ denote a vector of X that is at minimum distance from x (such a vector exists by the closure of X and Weierstrass' Theorem). By using the Lipschitz assumption (1.52) and the fact $c > L$, we have

$$F_c(x) = f(x) + c\|x - \hat{x}\| = f(\hat{x}) + \big(f(x) - f(\hat{x})\big) + c\|x - \hat{x}\| \ge f(\hat{x}) = F_c(\hat{x}),$$

with strict inequality if $x \ne \hat{x}$. Thus all minima of F_c over Y must lie in X, and also minimize f over X (since $F_c = f$ on X). Conversely, all minima

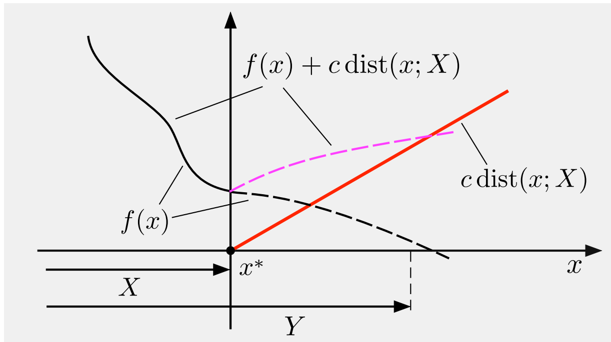

Figure 1.5.3. Illustration of Prop. 1.5.2. For c greater than the Lipschitz constant of f, the "slope" of the penalty function counteracts the "slope" of f at the optimal solution x^*.

of f over X are also minima of F_c over X (since $F_c = f$ on X), and by the preceding inequality, they are also minima of F_c over Y. **Q.E.D.**

The following proposition provides a generalization for constraints that involve the intersection of several sets.

Proposition 1.5.3: Let $f : \Re^n \mapsto \Re$ be a function, and let X_i, $i = 0, 1, \ldots, m$, be closed subsets of $\Re^n$ with nonempty intersection. Assume that f is Lipschitz continuous over X_0. Then there is a scalar $\overline{c} > 0$ such that for all $c \geq \overline{c}$, the set of minima of f over $\cap_{i=0}^m X_i$ coincides with the set of minima of

$$f(x) + c \sum_{i=1}^{m} \mathrm{dist}(x; X_i)$$

over X_0.

Proof: Let L be the Lipschitz constant for f, and let $c_1, \ldots, c_m$ be scalars satisfying

$$c_k > L + c_1 + \cdots + c_{k-1}, \qquad \forall\ k = 1, \ldots, m,$$

where $c_0 = 0$. Define

$$H_k(x) = f(x) + c_1\,\mathrm{dist}(x; X_1) + \cdots + c_k\,\mathrm{dist}(x; X_k), \qquad k = 1, \ldots, m,$$

and for $k = 0$, denote $H_0(x) = f(x)$, $c_0 = 0$. By applying Prop. 1.5.2, the set of minima of H_m over X_0 coincides with the set of minima of H_{m-1} over $X_m \cap X_0$, since c_m is greater than $L + c_1 + \cdots + c_{m-1}$, the Lipschitz constant for H_{m-1}. Similarly, for all $k = 1, \ldots, m$, the set of minima of

H_k over $\big(\cap_{i=k+1}^m X_i\big) \cap X_0$ coincides with the set of minima of H_{k-1} over $\big(\cap_{i=k}^m X_i\big) \cap X_0$. Thus, for $k = 1$, we obtain that the set of minima of H_m over X_0 coincides with the set of minima of H_0, which is f, over $\cap_{i=0}^m X_i$. Let

$$X^* \subset \cap_{i=0}^m X_i$$

be this set of minima. For $c \geq c_m$, we have $F_c \geq H_m$, while F_c coincides with H_m on X^*. Hence X^* is the set of minima of F_c over X_0. **Q.E.D.**

We finally note that exact penalty functions, and particularly the distance function $\text{dist}(x; X_i)$, are often relatively convenient in various contexts where difficult constraints complicate the algorithmic solution. As an example, see Section 6.4.4, where incremental proximal methods for highly constrained problems are discussed.

1.6 NOTES, SOURCES, AND EXERCISES

There is a very extensive literature on convex optimization, and in this section we will restrict ourselves to noting some books, research monographs, and surveys. In subsequent chapters, we will discuss in greater detail the literature that relates to the specialized content of these chapters.

Books relating primarily to duality theory are Rockafellar [Roc70], Stoer and Witzgall [StW70], Ekeland and Temam [EkT76], Bonnans and Shapiro [BoS00], Zalinescu [Zal02], Auslender and Teboulle [AuT03], and Bertsekas [Ber09].

The books by Rockafellar and Wets [RoW98], Borwein and Lewis [BoL00], and Bertsekas, Nedić, and Ozdaglar [BNO03] straddle the boundary between convex and variational analysis, a broad spectrum of topics that integrate classical analysis, convexity, and optimization of both convex and nonconvex (possibly nonsmooth) functions.

The book by Hiriart-Urruty and Lemarechal [HiL93] focuses on convex optimization algorithms. The books by Rockafellar [Roc84] and Bertsekas [Ber98] have a more specialized focus on network optimization algorithms and monotropic programming problems, which will be discussed in Chapters 4 and 6. The book by Ben-Tal and Nemirovski [BeN01] focuses on conic and semidefinite programming [see also the 2005 class notes by Nemirovski (on line), and the representative survey papers by Alizadeh and Goldfarb [AlG03], and Todd [Tod01]]. The book by Wolkowicz, Saigal, and Vanderberghe [WSV00] contains a collection of survey articles on semidefinite programming. The book by Boyd and Vanderberghe [BoV04] describes many applications, and contains a lot of related material and references. The book by Ben-Tal, El Ghaoui, and Nemirovski [BGN09] focuses on robust optimization; see also the survey by Bertsimas, Brown, and Caramanis [BBC11]. The book by Bauschke and Combettes [BaC11] develops the connection of convex analysis and monotone operator theory in infinite dimensional spaces. The book by Rockafellar and Wets [RoW98] also has a substantial finite-dimensional treatment of this subject. The books by Cottle, Pang, and Stone [CPS92], and Facchinei and Pang [FaP03] focus on complementarity and variational inequality problems. The books by Palomar and Eldar [PaE10], and Vetterli, Kovacevic, and Goyal [VKG14], and the surveys in the May 2010 issue of the IEEE Signal Processing Magazine describe applications of convex optimization in communications and signal processing. The books by Hastie, Tibshirani, and Friedman [HTF09], and Sra, Nowozin, and Wright [SNW12] describe applications of convex optimization in machine learning.

EXERCISES

1.1 (Support Vector Machines and Duality)

Consider the classification problem associated with a support vector machine,

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\|x\|^2 + \beta \sum_{i=1}^m \max\big\{0, 1 - b_i(c_i'x + y)\big\} \\ \text{subject to} & x \in \Re^n,\ y \in \Re, \end{array}$$

with quadratic regularization, where β is a positive regularization parameter (cf. Example 1.3.3).

(a) Write the problem in the equivalent form

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\|x\|^2 + \beta \sum_{i=1}^m \xi_i \\ \text{subject to} & x \in \Re^n,\ y \in \Re, \\ & 0 \le \xi_i, \quad 1 - b_i(c_i'x + y) \le \xi_i,\ i = 1, \ldots, m. \end{array}$$

Associate dual variables $\mu_i \ge 0$ with the constraints $1 - b_i(c_i'x + y) \le \xi_i$, and show that the dual function is given by

$$q(\mu) = \begin{cases} \hat{q}(\mu) & \text{if } \sum_{j=1}^m \mu_j b_j = 0,\ 0 \le \mu_i \le \beta,\ i = 1, \ldots, m, \\ -\infty & \text{otherwise,} \end{cases}$$

where

$$\hat{q}(\mu) = \sum_{i=1}^m \mu_i - \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m b_i b_j c_i' c_j \mu_i \mu_j.$$

Does the dual problem, viewed as the equivalent quadratic program

$$\begin{array}{ll} \text{minimize} & \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m b_i b_j c_i' c_j \mu_i \mu_j - \sum_{i=1}^m \mu_i \\ \text{subject to} & \displaystyle\sum_{j=1}^m \mu_j b_j = 0, \quad 0 \le \mu_i \le \beta,\ i = 1, \ldots, m, \end{array}$$

always have a solution? Is the solution unique? *Note*: The dual problem may have high dimension, but it has a generally more favorable structure than the primal. The reason is the simplicity of its constraint set, which makes it suitable for special types of quadratic programming methods, and the two-metric projection and coordinate descent methods of Section 2.1.2.

(b) Consider an alternative formulation where the variable y is set to 0, leading to the problem

$$\begin{aligned} &\text{minimize} \quad \tfrac{1}{2}\|x\|^2 + \beta \textstyle\sum_{i=1}^m \max\{0, 1 - b_i c_i' x\} \\ &\text{subject to} \ \ x \in \Re^n. \end{aligned}$$

Show that the dual problem should be modified so that the constraint $\sum_{j=1}^m \mu_j b_j = 0$ is not present, thus leading to a bound-constrained quadratic dual problem.

Note: The literature of the support vector machine field is extensive. Many of the nondifferentiable optimization methods to be discussed in subsequent chapters have been applied in connection to this field; see e.g., [MaM01], [FeM02], [SmS04], [Bot05], [Joa06], [JFY09], [JoY09], [SSS07], [LeW11].

1.2 (Minimizing the Sum or the Maximum of Norms [LVB98])

Consider the problems

$$\begin{aligned} &\text{minimize} \quad \sum_{i=1}^p \|F_i x + g_i\| \\ &\text{subject to} \ \ x \in \Re^n, \end{aligned} \tag{1.53}$$

and

$$\begin{aligned} &\text{minimize} \quad \max_{i=1,\ldots,p} \|F_i x + g_i\| \\ &\text{subject to} \ \ x \in \Re^n, \end{aligned}$$

where F_i and g_i are given matrices and vectors, respectively. Convert these problems to second order cone form and derive the corresponding dual problems.

1.3 (Complex l_1 and l_∞ Approximation [LVB98])

Consider the complex l_1 approximation problem

$$\begin{aligned} &\text{minimize} \quad \|Ax - b\|_1 \\ &\text{subject to} \ \ x \in \mathcal{C}^n, \end{aligned}$$

where $\mathcal{C}^n$ is the set of n-dimensional vectors whose components are complex numbers, and A and b are given matrix and vector with complex components. Show that it is a special case of problem (1.53) and derive the corresponding dual problem. Repeat for the complex l_∞ approximation problem

$$\begin{aligned} &\text{minimize} \quad \|Ax - b\|_\infty \\ &\text{subject to} \ \ x \in \mathcal{C}^n. \end{aligned}$$

1.4

The purpose of this exercise is to show that the SOCP can be viewed as a special case of SDP.

(a) Show that a vector $x \in \Re^n$ belongs to the second order cone if and only if the matrix

$$x_n I + \begin{pmatrix} 0 & 0 & \cdots & 0 & x_1 \\ 0 & 0 & \cdots & 0 & x_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & x_{n-1} \\ x_1 & x_2 & \cdots & x_{n-1} & 0 \end{pmatrix}$$

is positive semidefinite. *Hint*: We have that for any positive definite symmetric $n \times n$ matrix A, vector $b \in \Re^n$, and scalar d, the matrix

$$\begin{pmatrix} A & b \\ b' & c \end{pmatrix}$$

is positive definite if and only if

$$c - b'A^{-1}b > 0.$$

(b) Use part (a) to show that the primal SOCP can be written in the form of the dual SDP.

1.5 (Explicit Form of a Second Order Cone Problem)

Consider the SOCP (1.24).

(a) Partition the $n_i \times (n+1)$ matrices $(A_i \quad b_i)$ as

$$(A_i \quad b_i) = \begin{pmatrix} D_i & d_i \\ p_i' & q_i \end{pmatrix}, \qquad i = 1, \ldots, m,$$

where D_i is an $(n_i - 1) \times n$ matrix, $d_i \in \Re^{n_i - 1}$, $p_i \in \Re^n$, and $q_i \in \Re$. Show that

$$A_i x - b_i \in C_i \qquad \text{if and only if} \qquad \|D_i x - d_i\| \le p_i' x - q_i,$$

so we can write the SOCP (1.24) as

$$\begin{aligned} &\text{minimize} \quad c'x \\ &\text{subject to} \quad \|D_i x - d_i\| \le p_i' x - q_i, \ i = 1, \ldots, m. \end{aligned}$$

(b) Similarly partition λ_i as

$$\lambda_i = \begin{pmatrix} \mu_i \\ \nu_i \end{pmatrix}, \qquad i = 1, \ldots, m,$$

where $\mu_i \in \Re^{n_i-1}$ and $\nu_i \in \Re$. Show that the dual problem (1.25) can be written in the form

$$\begin{aligned} &\text{maximize} \quad \sum_{i=1}^{m} (d_i'\mu_i + q_i\nu_i) \\ &\text{subject to} \quad \sum_{i=1}^{m} (D_i'\mu_i + \nu_i p_i) = c, \quad \|\mu_i\| \le \nu_i,\ i = 1, \ldots, m. \end{aligned}$$

(c) Show that the primal and dual interior point conditions for strong duality (Prop. 1.2.4) hold if there exist primal and dual feasible solutions $\overline{x}$ and $(\overline{\mu}_i, \overline{\nu}_i)$ such that

$$\|D_i\overline{x} - d_i\| < p_i'\overline{x} - q_i, \qquad i = 1, \ldots, m,$$

and

$$\|\overline{\mu}_i\| < \overline{\nu}_i, \qquad i = 1, \ldots, m,$$

respectively.

1.6 (Separable Conic Problems)

Consider the problem

$$\begin{aligned} &\text{minimize} \quad \sum_{i=1}^{m} f_i(x_i) \\ &\text{subject to} \quad x \in S \cap C, \end{aligned}$$

where $x = (x_1, \ldots, x_m)$ with $x_i \in \Re^{n_i}$, $i = 1, \ldots, m$, and $f_i : \Re^{n_i} \mapsto (-\infty, \infty]$ is a proper convex function for each i, and S and C are a subspace and a cone of $\Re^{n_1+\cdots+n_m}$, respectively. Show that a dual problem is

$$\begin{aligned} &\text{maximize} \quad \sum_{i=1}^{m} q_i(\lambda_i) \\ &\text{subject to} \quad \lambda \in \hat{C} + S^{\perp}, \end{aligned}$$

where $\lambda = (\lambda_1, \ldots, \lambda_m)$, $\hat{C}$ is the dual cone of C, and

$$q_i(\lambda_i) = \inf_{z_i \in \Re} \big\{ f_i(z_i) - \lambda_i' z_i \big\}, \qquad i = 1, \ldots, m.$$

1.7 (Weber Points)

Consider the problem of finding a circle of minimum radius that contains r points $y_1, \ldots, y_r$ in the plane, i.e., find x and z that minimize z subject to $\|x - y_j\| \leq z$ for all $j = 1, \ldots, r$, where x is the center of the circle under optimization.

(a) Introduce multipliers μ_j, $j = 1, \ldots, r$, for the constraints, and show that the dual problem has an optimal solution and there is no duality gap.

(b) Show that calculating the dual function at some $\mu \geq 0$ involves the computation of a Weber point of $y_1, \ldots, y_r$ with weights $\mu_1, \ldots, \mu_r$, i.e., the solution of the problem

$$\min_{x \in \Re^2} \sum_{j=1}^{r} \mu_j \|x - y_j\|$$

(see Example 1.3.7).

1.8 (Inconsistent Convex Systems of Inequalities)

Let $g_j : \Re^n \mapsto \Re$, $j = 1, \ldots, r$, be convex functions over the nonempty convex set $X \subset \Re^n$. Show that the system

$$g_j(x) < 0, \qquad j = 1, \ldots, r,$$

has no solution within X if and only if there exists a vector $\mu \in \Re^r$ such that

$$\sum_{j=1}^{r} \mu_j = 1, \qquad \mu \geq 0,$$

$$\mu' g(x) \geq 0, \qquad \forall\, x \in X.$$

Note: This is an example of what is known as a *theorem of the alternative*. There are many results of this type, with a long history, such as the Farkas Lemma, and the theorems of Gordan, Motzkin, and Stiemke, which address the feasibility (possibly strict) of linear inequalities. They can be found in many sources, including Section 5.6 of [Ber09]. *Hint*: Consider the convex program

$$\begin{aligned} &\text{minimize } \ y \\ &\text{subject to } \ x \in X, \quad y \in \Re, \qquad g_j(x) \leq y, \quad j = 1, \ldots, r. \end{aligned}$$