

LECTURE 13: PRIMAL-DUAL OPTIMIZATION I

Index terms: nonsmooth problems, constrained optimization, duality, primal-dual methods

1 Introduction

In this lecture, we will delve into the realm of general nonsmooth problems and explore their underlying structure. We begin by introducing the formulation of composite minimization problems. These problems involve minimizing a sum of two convex and nonsmooth functions: $f(\mathbf{x})$ and $g(\mathbf{Ax})$, where \mathbf{A} is a linear operator. In particular, we focus on the following unconstrained optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) + g(\mathbf{Ax}). \quad (1)$$

Examples of such problems include scenarios where $g(\mathbf{Ax})$ could be the ℓ_1 or ℓ_2 -norm of the difference between \mathbf{Ax} and a vector \mathbf{b} ; i.e., $\|\mathbf{Ax} - \mathbf{b}\|_1$ and $\|\mathbf{Ax} - \mathbf{b}\|_2^2$. Alternatively, it can be the indicator function, denoted as $\delta_{\{\mathbf{b}\}}(\mathbf{Ax})$, which is zero when $\mathbf{Ax} = \mathbf{b}$ and infinity otherwise. The indicator function plays a crucial role in capturing constrained optimization problems. In this specific example, it corresponds to the following minimization $\min_{\mathbf{x} \in \mathcal{X}} \{f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b}\}$. Since solving a constrained problem is typically more challenging than solving an unconstrained problem, this composite formulation, with the assistance of indicator functions, may offer advantages. With the development of specialized algorithms for such problems, we can effectively address and solve these constrained optimization challenges. This would highlight the broad applicability and efficiency of composite problems, making it a valuable tool in optimization.

Next, we will employ a tool known as Fenchel conjugation to transform the formulation in (1) into a minimax problem. This concept is similar to what we discussed when addressing GANs (Generative Adversarial Networks), where we encountered a nonsmooth problem. By leveraging a technique called Kantorovich-Rubinstein duality, we were able to convert it into a minimax form with differentiable terms. This technique aligns with the principles we are about to explore through Fenchel conjugation.

2 Conjugation of functions

In optimization and convex analysis, the concept of conjugation of functions plays a fundamental role. It provides us with a powerful tool for characterizing and understanding the properties of convex functions. In this section, we delve into the definition and properties of Fenchel conjugation, a concept that enables us to explore the dual relationships between convex functions.

Definition 1. (Fenchel conjugation) Given a *proper, closed and convex function* $f : \mathcal{Q} \rightarrow \mathbb{R} \cup \{+\infty\}$, the function $f^* : \mathcal{Q}^* \rightarrow \mathbb{R} \cup \{+\infty\}$ such that

$$f^*(\mathbf{y}) = \sup_{x \in \text{dom}(f)} \{\mathbf{y}^T \mathbf{x} - f(x)\}$$

is called the *Fenchel conjugate* (or *conjugate*) of f .

When we evaluate the conjugate function at a fixed vector \mathbf{y} , it simply represents the maximum difference between the linear function $\mathbf{x}^T \mathbf{y}$ and the function $f(x)$. To gain a clearer geometric understanding, we can refer to Figure 1. In this visual representation, we can interpret \mathbf{y} as the slope. The maximum distance occurs at the point $\hat{\mathbf{x}}$. By lowering the red




Figure 1: The conjugate function $f^*(\mathbf{y})$ is the maximum gap between the linear function $\mathbf{x}^T \mathbf{y}$ (red line) and $f(\mathbf{x})$.

line by an amount equal to $\mathbf{y}^T \hat{\mathbf{x}} - f(\hat{\mathbf{x}})$, the shifted line (represented by the red dashed line) intersects the vertical axis at the point $-f^*(\mathbf{y})$. Thus, when we want to evaluate the conjugate function at a specific point \mathbf{y} , we consider all linear functions with a slope of \mathbf{y} and identify the one that lies below the function $f(\mathbf{x})$ (depicted by the blue curve) and just touches it. The point where this linear function intersects the vertical axis represents the value of the conjugate function. Now, let's delve into the crucial properties of conjugate function.

2.1 Properties of conjugate function

1. f^* is a convex and lower semicontinuous function by construction as the supremum of affine functions of \mathbf{y} .
2. The **conjugate** of the **conjugate** of a convex function f is the same function f ; i.e., $f^{**} = f$ for $f \in \mathcal{F}(Q)$.
3. The **conjugate** of the **conjugate** of a non-convex function f is its lower convex envelope when Q is compact:

$$f^{**}(\mathbf{x}) = \sup\{g(\mathbf{x}) : g \text{ is convex and } g \leq f, \forall \mathbf{x} \in Q\}.$$

Proof of Property 2 and 3. For any f , we have

$$\begin{aligned} f^{**}(\mathbf{x}) &= \sup_{\mathbf{y} \in \text{dom}(f^*)} \mathbf{x}^T \mathbf{y} - \left(\sup_{\hat{\mathbf{x}} \in \text{dom}(f)} \mathbf{y}^T \hat{\mathbf{x}} - f(\hat{\mathbf{x}}) \right) \\ &= \sup_{\mathbf{y} \in \text{dom}(f^*)} \mathbf{x}^T \mathbf{y} + \left(\inf_{\hat{\mathbf{x}} \in \text{dom}(f)} f(\hat{\mathbf{x}}) - \mathbf{y}^T \hat{\mathbf{x}} \right) \\ &= \sup_{\mathbf{y} \in \text{dom}(f^*), t \in \mathbb{R}} \mathbf{y}^T \mathbf{x} + t \\ &\quad \text{s.t. } \mathbf{y}^T \hat{\mathbf{x}} + t \leq f(\hat{\mathbf{x}}) \quad \forall \hat{\mathbf{x}} \in \text{dom}(f), \end{aligned}$$

where $\text{dom}(f) = Q$. Thus, f^{**} is the pointwise supremum of all affine functions not larger than f . By the envelope theorem, $f^{**} = \text{cl}(\text{conv}(f))$. If f is proper, convex and closed, then $f = \text{cl}(\text{conv}(f))$ and thus $f^{**} = f$. \square

4. [4, Theorem 3] For closed convex f , μ -strong convexity w.r.t. $\|\cdot\|$ is equivalent to $\frac{1}{\mu}$ smoothness of f^* w.r.t. $\|\cdot\|_*$.

Before proving proof of Property 4, we define strong smoothness to show a duality relationship between strong convexity and strong smoothness.

Definition 2. (Strong smoothness) A function $f : Q \rightarrow \mathbb{R}$ is μ -strongly smooth w.r.t. a norm $\|\cdot\|$ if f is everywhere differentiable and if for all \mathbf{x}, \mathbf{y} we have

$$f(\mathbf{x} + \mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} \rangle + \frac{1}{2} \mu \|\mathbf{y}\|^2.$$

Proof. First, [8, Lemma 15] yields one half of the claim (f strongly convex $\Rightarrow f^*$ strongly smooth). It is left to prove that f is strongly convex assuming that f^* is strongly smooth. For simplicity assume that $\mu = 1$. Denote $g(\mathbf{y}) = f^*(\mathbf{x} + \mathbf{y}) - (f^*(\mathbf{x}) + \langle \nabla f^*(\mathbf{x}), \mathbf{y} \rangle)$. By the smoothness assumption, $g(\mathbf{y}) \leq \frac{1}{2} \|\mathbf{y}\|_*^2$. This implies that $g^*(\mathbf{a}) \geq \frac{1}{2} \|\mathbf{a}\|^2$ because of [9, Lemma 19] and that the conjugate of half squared norm is half squared of the dual norm. Using the definition of g we have

$$\begin{aligned} g^*(\mathbf{a}) &= \sup_{\mathbf{y}} \langle \mathbf{y}, \mathbf{a} \rangle - g(\mathbf{y}) \\ &= \sup_{\mathbf{y}} \langle \mathbf{y}, \mathbf{a} \rangle - (f^*(\mathbf{x} + \mathbf{y}) - (f^*(\mathbf{x}) + \langle \nabla f^*(\mathbf{x}), \mathbf{y} \rangle)) \\ &= \sup_{\mathbf{y}} \langle \mathbf{y}, \mathbf{a} + \nabla f^*(\mathbf{x}) \rangle - f^*(\mathbf{x} + \mathbf{y}) + f^*(\mathbf{x}) \\ &= \sup_{\mathbf{z}} \langle \mathbf{z} - \mathbf{x}, \mathbf{a} + \nabla f^*(\mathbf{x}) \rangle - f^*(\mathbf{z}) + f^*(\mathbf{x}) \\ &= f(\mathbf{a} + \nabla f^*(\mathbf{x})) + f^*(\mathbf{x}) - \langle \mathbf{x}, \mathbf{a} + \nabla f^*(\mathbf{x}) \rangle \end{aligned}$$

where we have used that $f^{**} = f$, in the last step. Denote $\mathbf{u} = \nabla f^*(\mathbf{x})$. From the equality in Fenchel-Young (e.g. [9, Lemma 17]) we obtain that $\langle \mathbf{x}, \mathbf{u} \rangle = f^*(\mathbf{x}) + f(\mathbf{u})$ and thus

$$g^*(\mathbf{a}) = f(\mathbf{a} + \mathbf{u}) - f(\mathbf{u}) - \langle \mathbf{x}, \mathbf{a} \rangle.$$

Combining with $g^*(\mathbf{a}) \geq \frac{1}{2} \|\mathbf{a}\|^2$, we have

$$f(\mathbf{a} + \mathbf{u}) - f(\mathbf{u}) - \langle \mathbf{x}, \mathbf{a} \rangle \geq \frac{1}{2} \|\mathbf{a}\|^2, \quad (2)$$

which holds for all \mathbf{a}, \mathbf{x} , with $\mathbf{u} = \nabla f^*(\mathbf{x})$.

Now let us prove that for any point \mathbf{u}' in the relative interior of the domain of f that if $\mathbf{x} \in \partial f(\mathbf{u}')$ then $\mathbf{u}' = \nabla f^*(\mathbf{x})$. Let $\mathbf{u} := \nabla f^*(\mathbf{x})$ and we must show that $\mathbf{u}' = \mathbf{u}$. By Fenchel-Young, we have that $\langle \mathbf{x}, \mathbf{u}' \rangle = f^*(\mathbf{x}) + f(\mathbf{u}')$, and, again by Fenchel-Young (and $f^{**} = f$), we have $\langle \mathbf{x}, \mathbf{u} \rangle = f^*(\mathbf{x}) + f(\mathbf{u})$. We can now apply Equation (2), to obtain:

$$\begin{aligned} 0 &= \langle \mathbf{x}, \mathbf{u} \rangle - f(\mathbf{u}) - (\langle \mathbf{x}, \mathbf{u}' \rangle - f(\mathbf{u}')) \\ &= f(\mathbf{u}') - f(\mathbf{u}) - \langle \mathbf{x}, \mathbf{u}' - \mathbf{u} \rangle \geq \frac{1}{2} \|\mathbf{u}' - \mathbf{u}\|^2, \end{aligned}$$

which implies that $\mathbf{u}' = \nabla f^*(\mathbf{x})$.

Next, let $\mathbf{u}_1, \mathbf{u}_2$ be two points in the relative interior of the domain of f , let $\alpha \in (0, 1)$, and let $\mathbf{u} = \alpha \mathbf{u}_1 + (1 - \alpha) \mathbf{u}_2$. Let $\mathbf{x} \in \partial f(\mathbf{u})$ (which is non-empty¹). We have that $\mathbf{u} = \nabla f^*(\mathbf{x})$, by the previous argument. Now we are able to apply Equation (2) twice, once with $\mathbf{a} = \mathbf{u}_1 - \mathbf{u}$ and once with $\mathbf{a} = \mathbf{u}_2 - \mathbf{u}$ (and both with \mathbf{x}) to obtain

$$\begin{aligned} f(\mathbf{u}_1) - f(\mathbf{u}) - \langle \mathbf{x}, \mathbf{u}_1 - \mathbf{u} \rangle &\geq \frac{1}{2} \|\mathbf{u}_1 - \mathbf{u}\|^2, \\ f(\mathbf{u}_2) - f(\mathbf{u}) - \langle \mathbf{x}, \mathbf{u}_2 - \mathbf{u} \rangle &\geq \frac{1}{2} \|\mathbf{u}_2 - \mathbf{u}\|^2. \end{aligned}$$

Finally, summing up the above two equations with coefficients α and $1 - \alpha$ we obtain that f is strongly convex. \square

Let's now exemplify the conjugate of some commonly used functions.

Example 1. (Conjugate of ℓ_2 -norm-squared) Let $f(\mathbf{x}) = \frac{\lambda}{2} \|\mathbf{x}\|^2$, then f is λ -strongly convex and λ -smooth. From the Property 4, we can expect the dual to be $\frac{1}{\lambda}$ -smooth and -strongly convex. By definition $f^*(\mathbf{y}) = \max_{\mathbf{x}} \langle \mathbf{y}, \mathbf{x} \rangle - \frac{\lambda}{2} \|\mathbf{x}\|^2$, which is concave quadratic in \mathbf{x} , and we can find the maximum by setting the gradient equal to 0.

$$0 = \mathbf{y} - \lambda \mathbf{x}^* \iff \mathbf{x}^* = \frac{1}{\lambda} \mathbf{y} \iff f^*(\mathbf{y}) = \frac{1}{\lambda} \|\mathbf{y}\|^2 - \frac{1}{2\lambda} \|\mathbf{y}\|^2 = \frac{1}{2\lambda} \|\mathbf{y}\|^2.$$

It is $\frac{1}{\lambda}$ -smooth and -strongly convex, as expected.

¹The set $\partial f(\mathbf{u})$ is not empty for all \mathbf{u} in the relative interior of the domain of f . See the relative max formula in [2, page 42] or [7, page 253]. If \mathbf{u} is not in the interior of f , then $\partial f(\mathbf{u})$ is empty. But, a function is defined to be essentially strictly convex if it is strictly convex on any subset of $\{\mathbf{u} : \partial f(\mathbf{u}) \neq \emptyset\}$. The last set is called the domain of ∂f and it contains the relative interior of the domain of f , so we are fine here.

Example 2. (Conjugate of ℓ_1 -norm) $f(\mathbf{x}) = \lambda\|\mathbf{x}\|_1 \Rightarrow f^*(\mathbf{y}) = \max_{\mathbf{x}} \langle \mathbf{y}, \mathbf{x} \rangle - \lambda\|\mathbf{x}\|_1$. By definition of the ℓ_1 -norm:

$$f^*(\mathbf{y}) = \max_{\mathbf{x}} \sum_{i=1}^n y_i x_i - \lambda|x_i| = \max_{\mathbf{x}} \sum_{i=1}^n y_i \text{sign}(x_i) |x_i| - \lambda|x_i| = \sum_{i=1}^n \max_{x_i} y_i \text{sign}(x_i) |x_i| - \lambda|x_i|.$$

The second equality decomposes x_i into its sign and its absolute value. The third decomposes the maximization over the components. Since this is a non-differentiable at 0, let's observe it case by case.

- If all $|y_i| \leq \lambda$, then $\forall i, (y_i \text{sign}(x_i) - \lambda)|x_i| \leq 0$. Taking $\mathbf{x} = 0$ gives the maximum value: $f^*(\mathbf{y}) = 0$.
- If for at least one $i, |y_i| > \lambda$, $(y_i \text{sign}(x_i) - \lambda)|x_i| \rightarrow +\infty$ as $|x_i| \rightarrow +\infty$.

Therefore, $f^*(\mathbf{y}) = \delta_{\mathbf{y}: \|\mathbf{y}\|_\infty \leq \lambda}(\mathbf{y}) = \begin{cases} 0, & \text{if } \|\mathbf{y}\|_\infty \leq \lambda, \\ +\infty, & \text{if } \|\mathbf{y}\|_\infty > \lambda. \end{cases}$

Recall from the dual norm $\lambda\|\mathbf{y}\|_1 = \sup_{\mathbf{x}} \{\langle \mathbf{x}, \mathbf{y} \rangle : \|\mathbf{x}\|_\infty \leq \lambda\}$. Additionally, as discussed in Section 1, we can convert a constrained optimization problem into an unconstrained one using an indicator function. When we perform this conversion, we have the following equivalence:

$$\lambda\|\mathbf{y}\|_1 = \sup_{\mathbf{x}} \{\langle \mathbf{x}, \mathbf{y} \rangle : \|\mathbf{x}\|_\infty \leq \lambda\} = \sup_{\mathbf{x}} \{\langle \mathbf{x}, \mathbf{y} \rangle - \delta_{\mathbf{y}: \|\mathbf{y}\|_\infty \leq \lambda}(\mathbf{y})\} = \delta_{\mathbf{y}: \|\mathbf{y}\|_\infty \leq \lambda}^*(\mathbf{y}),$$

where $\delta_{\mathbf{y}: \|\mathbf{y}\|_\infty \leq \lambda}^*$ is the conjugate of $\delta_{\mathbf{y}: \|\mathbf{y}\|_\infty \leq \lambda}$. This observation also illustrates a specific instance of a more general principle: the Fenchel conjugate of a norm corresponds to the indicator function defined on the unit ball of the dual norm. In this case, we observe this relationship for the ℓ_1 -norm and the ℓ_∞ -norm.

3 General non-smooth problems

Let's go back to our composite formulation.

$$\min_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) + g(\mathbf{Ax})$$

Now by using Fenchel-conjugation of a convex function and its property in Property 2, $g(\mathbf{Ax}) = \max_{\mathbf{y}} \langle \mathbf{Ax}, \mathbf{y} \rangle - g^*(\mathbf{y})$, where g^* is the conjugate of g , we can reformulate it as a **minimax** problem.

$$\min_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) + g(\mathbf{Ax}) = \min_{\mathbf{x} \in \mathbb{R}^p} \max_{\mathbf{y}} \{\Phi(\mathbf{x}, \mathbf{y}) := f(\mathbf{x}) + \langle \mathbf{Ax}, \mathbf{y} \rangle - g^*(\mathbf{y})\}$$

Because $\langle \mathbf{Ax}, \mathbf{y} \rangle - g^*(\mathbf{y})$ is a sum of an affine and a concave function, it is concave in \mathbf{y} . $f(\mathbf{x})$ is already assumed to be convex in \mathbf{x} , hence we obtain a **convex-concave min-max** problem. As we will see in Lecture 14, there are efficient methods to solve such problems such as the Extra-Gradient algorithm [5].

3.1 Min-max reformulation

We now see a special case of (1), i.e., when g is the indicator function to address the linearly constrained problems. Suppose we want to solve

$$\min_{\mathbf{x}} \{f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b}\}. \quad (3)$$

To this end, we first reformulate it as an unconstrained problem via indicator function trick, then we replace the indicator function with its conjugate formulation to obtain a min-max problem. In this case,

$$g(\mathbf{Ax}) = \delta_{\{\mathbf{b}\}}(\mathbf{Ax}) = \begin{cases} 0, & \text{if } \mathbf{Ax} = \mathbf{b}, \\ +\infty, & \text{if } \mathbf{Ax} \neq \mathbf{b}. \end{cases}$$

$$g^*(\mathbf{y}) = \max_{\mathbf{x}} \langle \mathbf{y}, \mathbf{x} \rangle - \delta_{\{\mathbf{b}\}}(\mathbf{x}) = \max_{\mathbf{x}: \mathbf{x}=\mathbf{b}} \langle \mathbf{y}, \mathbf{x} \rangle = \langle \mathbf{y}, \mathbf{b} \rangle.$$

The second identity holds because if $\mathbf{x} \neq \mathbf{b}$, negative indicator function evaluates to $-\infty$. However, since the problem maximizes it would push \mathbf{x} to be equal to \mathbf{b} .

We reach the minimax formulation via conjugation:

$$\min_{\mathbf{x}} \{f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b}\} = \min_{\mathbf{x}} f(\mathbf{x}) + g(\mathbf{Ax}) = \min_{\mathbf{x}} \max_{\mathbf{y}} f(\mathbf{x}) + \langle \mathbf{Ax} - \mathbf{b}, \mathbf{y} \rangle. \quad (4)$$

This is what we call the ‘‘Lagrangian’’ formulation. It is particularly intriguing because when we set up this min-max problem, we are essentially making sure that the constraint $\mathbf{Ax} = \mathbf{b}$ is satisfied; otherwise, since $\langle \mathbf{Ax} - \mathbf{b}, \mathbf{y} \rangle$ is a linear function in \mathbf{y} , it can easily go to infinity. If we consider this as a game between two players with playing strategies \mathbf{x} and \mathbf{y} , the ‘‘ \mathbf{y} ’’ player forces the ‘‘ \mathbf{x} ’’ player to satisfy $\mathbf{Ax} = \mathbf{b}$ while ‘‘ \mathbf{x} ’’ player is still trying to minimize the function $f(\mathbf{x})$. This framework can be considered as the [bilinear min-max template](#). More generally, the bilinear min-max template has the following form.

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}) + \langle \mathbf{Ax}, \mathbf{y} \rangle - h(\mathbf{y}),$$

where $\mathcal{X} \subseteq \mathbb{R}^p$ and $\mathcal{Y} \subseteq \mathbb{R}^n$, $f: \mathcal{X} \rightarrow \mathbb{R}$ is convex, and $h: \mathcal{Y} \rightarrow \mathbb{R}$ is convex.

In the next example, we see that a potentially more complex problem can be formulated as in the template (3).

Example 3. (Sparse Recovery) In the context of MRI and sparse recovery, we use the Basis Pursuit Denoising (BPDN) formulation to find a sparse solution for an image. This formulation is expressed as:

$$\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{R}^p} \{ \|\mathbf{x}\|_1 : \|\mathbf{Ax} - \mathbf{b}\|_2 \leq \|\mathbf{w}\|_2, \|\mathbf{x}\|_\infty \leq 1 \}. \quad (\text{BPDN})$$

It seeks to find a sparse solution \mathbf{x} , representing the image we are recovering while satisfying two constraints: $\mathbf{Ax} - \mathbf{b}$ should be within an acceptable level of noise, where \mathbf{A} represents the sub-sampled Fourier data, and the image coefficients are bounded.

An alternative formulation:

An equivalent formulation to (BPDN) can be given as

$$f^* := \min_{\mathbf{x} \in \mathbb{R}^p} \{f(\mathbf{x}) : \mathbf{Ax} - \mathbf{b} \in \mathcal{K}, \mathbf{x} \in \mathcal{X}\}. \quad (5)$$

The above template captures BPDN formulation with

- $f(\mathbf{x}) = \|\mathbf{x}\|_1$, a convex function.
- $\mathcal{K} = \{\|\mathbf{u}\| \in \mathbb{R}^n : \|\mathbf{u}\| \leq \|\mathbf{w}\|_2\}$, a convex set.
- $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^p : \|\mathbf{x}\|_\infty \leq 1\}$, also a convex set.

Now that we simplified the formulation, as a next step, we will see how to reformulate this as the simple form:

$$\min_{\mathbf{x}} \{f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b}\}.$$

To this end, we will use a technique called *lifting* from convex optimization to address the problem by moving it into a higher-dimensional space.

Reformulation between templates:

As the first step, let us define some auxiliary variables $\mathbf{r}_1 = \mathbf{Ax} - \mathbf{b} \in \mathbb{R}^n$ and $\mathbf{r}_2 = \mathbf{x} \in \mathbb{R}^p$. Thus, solving the below problem is equivalent to solving (5).

$$\min_{\mathbf{x}, \mathbf{r}_1, \mathbf{r}_2} \{f(\mathbf{x}) : \mathbf{r}_1 \in \mathcal{K}, \mathbf{r}_2 \in \mathcal{X}, \mathbf{Ax} - \mathbf{b} = \mathbf{r}_1, \mathbf{x} = \mathbf{r}_2\}.$$

As the second step, let’s define $\mathbf{z} = \begin{bmatrix} \mathbf{x} \\ \mathbf{r}_1 \\ \mathbf{r}_2 \end{bmatrix} \in \mathbb{R}^{2p+n}$, $\bar{\mathbf{A}} = \begin{bmatrix} \mathbf{A} & -\mathbf{I}_{n \times n} & \mathbf{0}_{n \times p} \\ \mathbf{I}_{p \times p} & \mathbf{0}_{p \times n} & -\mathbf{I}_{p \times p} \end{bmatrix}$, $\bar{\mathbf{b}} = \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix}$, $\bar{f}(\mathbf{z}) = f(\mathbf{x}) + \delta_{\mathcal{K}}(\mathbf{r}_1) + \delta_{\mathcal{X}}(\mathbf{r}_2)$,

where $\delta_{\mathcal{X}}(\mathbf{x}) = 0$, if $\mathbf{x} \in \mathcal{X}$, and $\delta_{\mathcal{X}}(\mathbf{x}) = +\infty$, otherwise. The simplified template yields the lifted problem:

$$\min_{\mathbf{z} \in \mathbb{R}^{2p+n}} \{\bar{f}(\mathbf{z}) : \bar{\mathbf{A}}\mathbf{z} = \bar{\mathbf{b}}\}.$$

As a result, (BPDN), a potentially more complex problem, can be reformulated into a simpler convex function minimization problem with an affine constraint, as seen in (3). This transformation can offer significant advantages, especially if we have access to efficient algorithms to solve (3). However, it is worth noting that the process of dimension lifting can also introduce additional complexity. Since both (BPDN) and (3) represent convex optimization problems, the choice of formulation should be made carefully, taking into account the specific characteristics and requirements of the problem at hand.

There are other reformulation techniques for our general template (3), which can be found in the **standard convex optimization** literature such as *linear programming*, *convex quadratic programming*, *second order cone programming*, *semidefinite programming* and *geometric programming*. Since these formulations are well-documented, one can often leverage general-purpose solvers to tackle the problem. However, the feasibility and effectiveness of these techniques depend on the specific characteristics of the problem. Moreover, there is a plethora of other reformulation techniques available, centered around the idea of converting unconstrained problems into convex forms through techniques like **convex splitting**. Examples include *composite convex minimization* and *consensus optimization*, among others.

Going back to our minimax formulation, it is crucial to question whether we can swap the order of min and max within the minimax formulation and how it can be useful in finding a solution for such problems. The next section introduces answers to these questions through the duality concept.

4 Dual problem

In this section, we discuss when we can change the order of min and max in the minimax formulation. This interchange can yield certain advantages in solving the problem. To better understand this, we introduce the concept of the dual maximization problem:

$$\max_{\mathbf{y} \in \mathbb{R}^n} d(\mathbf{y}) := \max_{\mathbf{y} \in \mathbb{R}^n} \underbrace{\{\min_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) + \langle \mathbf{y}, \mathbf{A}\mathbf{x} - \mathbf{b} \rangle\}}_{d(\mathbf{y})}.$$

Note that even if $f(\mathbf{x})$ is not convex, $d(\mathbf{y})$ is concave. This is because for each \mathbf{x} , $d(\mathbf{y})$ is linear; i.e., it is both convex and concave. Furthermore, pointwise minimum of concave functions is still concave.

Remark: If we can exchange min and max, we obtain a **concave** maximization problem, which is 'usually' easy to deal with.

Example 4. Suppose $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$. Then $d(\mathbf{y}) := \max_{\mathbf{y} \in \mathbb{R}^n} \{\min_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{2}\|\mathbf{x}\|_2^2 - \langle \mathbf{y}, \mathbf{A}\mathbf{x} - \mathbf{b} \rangle\}$, and we are solving

$$\max_{\mathbf{y} \in \mathbb{R}^n} d(\mathbf{y}) = \max_{\mathbf{y} \in \mathbb{R}^n} \{\min_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{2}\|\mathbf{x}\|_2^2 - \langle \mathbf{y}, \mathbf{A}\mathbf{x} - \mathbf{b} \rangle\}.$$

Inner problem is a minimization of a quadratic convex function, hence, the minimum can be found by setting the gradient equal to 0.

$$\nabla_{\mathbf{x}} \left(\frac{1}{2}\|\mathbf{x}\|_2^2 - \langle \mathbf{y}, \mathbf{A}\mathbf{x} - \mathbf{b} \rangle \right) = \mathbf{x}^* - \mathbf{A}^\top \mathbf{y} = 0 \iff \mathbf{x}^* = \mathbf{A}^\top \mathbf{y} \iff d(\mathbf{y}) = \frac{1}{2}\mathbf{y}^\top \mathbf{A}\mathbf{A}^\top \mathbf{y} - \mathbf{y}^\top (\mathbf{A}\mathbf{A}^\top \mathbf{y} - \mathbf{b}) = -\frac{1}{2}\mathbf{y}^\top \mathbf{A}\mathbf{A}^\top \mathbf{y} + \mathbf{y}^\top \mathbf{b}$$

Then the min-max problem reduces to

$$\max_{\mathbf{y} \in \mathbb{R}^n} d(\mathbf{y}) = \max_{\mathbf{y} \in \mathbb{R}^n} -\frac{1}{2}\mathbf{y}^\top \mathbf{A}\mathbf{A}^\top \mathbf{y} + \mathbf{y}^\top \mathbf{b}$$

The objective is to maximize a negative quadratic function, thus again similarly global maximizer can be found by

$$\nabla_{\mathbf{y}}(d(\mathbf{y})) = -\mathbf{A}\mathbf{A}^\top \mathbf{y}^* + \mathbf{b} = 0 \iff \mathbf{y}^* = (\mathbf{A}\mathbf{A}^\top)^{-1} \mathbf{b}$$

Plugging this into \mathbf{x}^* , we get $\mathbf{x}^* = \mathbf{A}^\top (\mathbf{A}\mathbf{A}^\top)^{-1} \mathbf{b}$. This is the same solution as the least-squares solution!

In this example, we easily solved a minimax problem, although such problems are generally difficult. If certain conditions are met, we can apply the same techniques to solve other challenging problems. We will explore these conditions after discussing another example of a dual problem.

Example 5. (Non-smoothness of the dual function) Consider a constrained convex problem:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^3} \quad & \{f(\mathbf{x}) := x_1^2 + 2x_2\}, \\ \text{s.t.} \quad & 2x_3 - x_1 - x_2 = 1, \\ & \mathbf{x} \in \mathcal{X} := [-2, 2] \times [-2, 2] \times [0, 2]. \end{aligned}$$

Let $\lambda \in \mathbb{R}$ be the dual variable (or Lagrange multiplier) associated with the constraint $2x_3 - x_1 - x_2 = 1$. Notice that the number of constraints and the dimension of the dual variable should match (1 in this case). The **dual function**, as written below, is **concave**.

$$d(\lambda) := \min_{\mathbf{x} \in \mathcal{X}} \{x_1^2 + 2x_2 + \lambda(2x_3 - x_1 - x_2 - 1)\}.$$

The dual function's nature is further illustrated in Figure 5, with a particular emphasis on its **nonsmooth** characteristics. It is worth noting that, in general, dual functions tend to be non-smooth.




Figure 2: The plot of the dual function $d(\lambda)$ with respect to λ

4.1 Exchanging min and max order

In this section, we explore the fundamental principles of duality theory: *weak duality* and *strong duality*.

4.1.1 Weak duality

The value of the primal problem (minimization) is **always** greater than or equal to the value of its dual problem (maximization). This property is called weak duality. In other words,

$$\underbrace{\max_{\mathbf{y} \in \mathbb{R}^n} d(\mathbf{y})}_{\text{Dual problem}} =: \boxed{\max_{\mathbf{y} \in \mathbb{R}^n} \min_{\mathbf{x} \in \mathbb{R}^p} \Phi(\mathbf{x}, \mathbf{y}) \leq \min_{\mathbf{x} \in \mathbb{R}^p} \max_{\mathbf{y} \in \mathbb{R}^n} \Phi(\mathbf{x}, \mathbf{y})} = \underbrace{\min_{\mathbf{x} \in \mathbb{R}^p} \{f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b}\}}_{\text{Primal problem}} = \begin{cases} f^*, & \text{if } \mathbf{Ax} = \mathbf{b}, \\ +\infty, & \text{if } \mathbf{Ax} \neq \mathbf{b}. \end{cases}$$

A simple example of weak duality is shown in Figure 3.

Proof of weak duality. Let x^* and f^* be the solution to the primal problem and its corresponding optimal value, respectively. The primal is given below.

$$f^* := \min_{\mathbf{x} \in \mathbb{R}^p} \{f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b}\} = \min_{\mathbf{x} \in \mathbb{R}^p} \max_{\mathbf{y} \in \mathbb{R}^n} \{\Phi(\mathbf{x}, \mathbf{y}) := f(\mathbf{x}) + \langle \mathbf{y}, \mathbf{Ax} - \mathbf{b} \rangle\}$$

Since $\mathbf{Ax}^* = \mathbf{b}$, it holds for any \mathbf{y} ,

$$\begin{aligned} \Phi(\mathbf{x}^*, \mathbf{y}) &= f^* = f(\mathbf{x}^*) + \langle \mathbf{y}, \mathbf{Ax}^* - \mathbf{b} \rangle \\ &\geq \min_{\mathbf{x} \in \mathbb{R}^p} \{f(\mathbf{x}) + \langle \mathbf{y}, \mathbf{Ax} - \mathbf{b} \rangle\} \\ &= \min_{\mathbf{x} \in \mathbb{R}^p} \Phi(\mathbf{x}, \mathbf{y}). \end{aligned}$$




Figure 3: Visualization of weak duality in a specific example

Take maximum of both sides in \mathbf{y} and note that f^* is independent of \mathbf{y} :

$$f^* = \min_{\mathbf{x} \in \mathbb{R}^p} \max_{\mathbf{y} \in \mathbb{R}^n} \Phi(\mathbf{x}, \mathbf{y}) \geq \max_{\mathbf{y} \in \mathbb{R}^n} \min_{\mathbf{x} \in \mathbb{R}^p} \Phi(\mathbf{x}, \mathbf{y}) =: \max_{\mathbf{y} \in \mathbb{R}^n} d(\mathbf{y}) = d^*.$$

□

Now, let us analyze the case where the inequality in weak duality is tight; i.e., the concept of strong duality.

4.1.2 Strong duality

To understand the concept of strong duality, let's first recall the definition of a saddle point.

Definition 3. A point $(\mathbf{x}^*, \mathbf{y}^*) \in \mathbb{R}^p \times \mathbb{R}^n$ is called a *saddle point* of Φ if

$$\Phi(\mathbf{x}^*, \mathbf{y}) \leq \Phi(\mathbf{x}^*, \mathbf{y}^*) \leq \Phi(\mathbf{x}, \mathbf{y}^*), \quad \forall \mathbf{x} \in \mathbb{R}^p, \mathbf{y} \in \mathbb{R}^n.$$

To provide some visual intuition, Figure 4 illustrates what a saddle point looks like.




Figure 4: An illustration of saddle point.

Let us assume that there exists \mathbf{x}^* and $\Phi(\mathbf{x}, \mathbf{y})$ has a saddle point to avoid pathological situations. If the optimal values of the primal and its dual problem coincide, we say that strong duality holds. Namely, it is the case when we have

$$f^* = f(\mathbf{x}^*) = \min_{\mathbf{x} \in \mathbb{R}^p} \max_{\mathbf{y} \in \mathbb{R}^n} \Phi(\mathbf{x}, \mathbf{y}) = \max_{\mathbf{y} \in \mathbb{R}^n} \min_{\mathbf{x} \in \mathbb{R}^p} \Phi(\mathbf{x}, \mathbf{y}) =: \max_{\mathbf{y} \in \mathbb{R}^n} d(\mathbf{y}) = d^*.$$

The below example shows a case where strong duality holds.

Example 6. Consider the following primal minimization problem: $\min_{\mathbf{x}} P(\mathbf{x}) := f(\mathbf{x}) + g(\mathbf{x}) := \frac{1}{2}\|\mathbf{x}\|^2 + \|\mathbf{x}\|_1$ and suppose we know that strong duality holds². Then, using conjugation and strong duality

$$\begin{aligned} P(\mathbf{x}^*) &= \min_{\mathbf{x}} P(\mathbf{x}) = \min_{\mathbf{x}} \max_{\mathbf{y}} f(\mathbf{x}) + \langle \mathbf{x}, \mathbf{y} \rangle - g^*(\mathbf{y}), && \text{by conjugation} \\ &= \max_{\mathbf{y}} -g^*(\mathbf{y}) + \min_{\mathbf{x}} f(\mathbf{x}) + \langle \mathbf{x}, \mathbf{y} \rangle, && \text{by changing min-max} \\ &= \max_{\mathbf{y}} -g^*(\mathbf{y}) - \max_{\mathbf{x}} \langle \mathbf{x}, -\mathbf{y} \rangle - f(\mathbf{x}), && \text{by } \min f = -\max -f \\ &= \max_{\mathbf{y}} -g^*(\mathbf{y}) - f^*(-\mathbf{y}), && \text{by conjugation.} \end{aligned}$$

Thus, we can write the dual problem as $d^* = \max_{\mathbf{y}} d(\mathbf{y}) = -g^*(\mathbf{y}) - f^*(-\mathbf{y})$. Recall that in Example 1 and Example 2, we found the conjugates of ℓ_2 and ℓ_1 -norm to be $f^*(-\mathbf{y}) = \frac{1}{2}\|\mathbf{y}\|^2$ and $g^*(\mathbf{y}) = \delta_{\mathbf{y}: \|\mathbf{y}\|_\infty \leq 1}(\mathbf{y})$. Plugging these in yields

$$\text{Primal problem: } \min_{\mathbf{x}} P(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|^2 + \|\mathbf{x}\|_1$$

$$\text{Dual problem: } \max_{\mathbf{y}} -\frac{1}{2}\|\mathbf{y}\|^2 - \delta_{\mathbf{y}: \|\mathbf{y}\|_\infty \leq 1}(\mathbf{y})$$

We can also plot the primal and the dual to visualize the strong duality as in Figure 5. In addition, since our primal is strongly convex, the dual is expected to be Lipschitz smooth, as evident from the figure.




Figure 5: Primal (in black) and the dual (in red) functions showing the primal minimum and dual maximum coincide at $(\mathbf{x}, \mathbf{y}) = (0, 0)$.

Necessary and sufficient condition for strong duality

Note that the existence of a saddle point is not automatic even in a convex-concave setting. However, we see the conditions that guarantee strong duality. To this end, let's define the critical condition that plays a key role:

Definition 4. For the primal problem, *Slater's condition* is said to hold if $\mathcal{D} := \text{relint}(\text{dom} f) \cap \{\mathbf{x} : \mathbf{Ax} = \mathbf{b}\} \neq \emptyset$.

Remark. To provide a more intuitive understanding, we can also explain Slater's condition as the following.

²One needs to check if this before applying strong duality result. We will see the conditions under which strong duality holds after this example

- If $\text{dom} f = \mathbb{R}^p$, then Slater's condition holds $\Leftrightarrow \boxed{\exists \bar{\mathbf{x}} : \mathbf{A}\bar{\mathbf{x}} = \mathbf{b}}$.
- If $\text{dom} f = \mathbb{R}^p$ and instead of $\mathbf{Ax} = \mathbf{b}$, we have the feasible set $\{\mathbf{x} : h(\mathbf{x}) \leq 0\}$, where h is $\mathbb{R}^p \rightarrow \mathbb{R}^q$ is convex, then

$$\text{Slater's condition holds} \Leftrightarrow \boxed{\exists \bar{\mathbf{x}} : h(\bar{\mathbf{x}}) < 0.}$$

Now, we are ready to state the necessary and sufficient condition for strong duality.

Theorem 4.1. [3, Section 5.3.2](Necessary and sufficient optimality condition) *If the primal and dual problems are given by*

$$f^* := \begin{cases} \min_{\mathbf{x} \in \mathbb{R}^p} & f(\mathbf{x}) \\ \text{s.t.} & \mathbf{Ax} = \mathbf{b}, \end{cases} \quad \text{and} \quad d^* := \max_{\mathbf{y} \in \mathbb{R}^n} d(\mathbf{y}).$$

and Slater's condition is satisfied, then strong duality holds.

This is a very powerful theorem and it will prove valuable as it enables us to interchange the order of max and min!

Proof of Theorem 4.1. We consider the primal problem with f convex, and assume Slater's condition holds: There exists $\bar{\mathbf{x}} \in \text{relint}(\text{dom} f)$ and $\mathbf{A}\bar{\mathbf{x}} = \mathbf{b}$. In order to simplify the proof, we make two additional assumptions: first that $\text{dom} f$ has nonempty interior (hence, $\text{relint}(\text{dom} f) = \text{int}(\text{dom} f)$) and second, that $\text{rank}(\mathbf{A}) = n$ ($\mathbf{A} \in \mathbb{R}^{n \times p}$). We assume that f^* is finite. (Since there is a feasible point, we can only have $f^* = -\infty$ or f^* finite; if $f^* = -\infty$, then $d^* = -\infty$ by weak duality.) We begin by establishing a simple geometric interpretation of the dual function in terms of some sets. First, consider the set \mathcal{G} defined as:

$$\mathcal{G} := \{(\mathbf{Ax} - \mathbf{b}, f(\mathbf{x})) \in \mathbb{R}^n \times \mathbb{R} \mid \mathbf{x} \in \text{dom} f\},$$

which is the set of values taken on by the constraint and objective functions. The optimal value f^* is easily expressed in terms of \mathcal{G} as

$$f^* = \inf\{t \mid (\mathbf{v}, t) \in \mathcal{G}, \mathbf{v} = 0\}.$$

Let $\boldsymbol{\lambda} \in \mathbb{R}^n$ be the dual variable associated with the constraint $\mathbf{Ax} = \mathbf{b}$. To evaluate the dual function at $\boldsymbol{\lambda}$, we minimize the affine function

$$(\boldsymbol{\lambda}, 1)^T(\mathbf{v}, t) = \sum_{i=1}^n \lambda_i v_i + t$$

over $(\mathbf{v}, t) \in \mathcal{G}$, i.e., we have

$$g(\boldsymbol{\lambda}) = \inf\{(\boldsymbol{\lambda}, 1)^T(\mathbf{v}, t) \mid (\mathbf{v}, t) \in \mathcal{G}\}.$$

In particular, we see that if the infimum is finite, then the inequality

$$(\boldsymbol{\lambda}, 1)^T(\mathbf{v}, t) \geq g(\boldsymbol{\lambda})$$

defines a supporting hyperplane to \mathcal{G} . This is sometimes referred to as a nonvertical supporting hyperplane because the last component of the normal vector is nonzero. Obviously, $t \geq (\boldsymbol{\lambda}, 1)^T(\mathbf{v}, t)$ if $\mathbf{v} = 0$. Therefore

$$\begin{aligned} f^* &= \inf\{t \mid (\mathbf{v}, t) \in \mathcal{G}, \mathbf{v} = 0\} \\ &\geq \inf\{(\boldsymbol{\lambda}, 1)^T(\mathbf{v}, t) \mid (\mathbf{v}, t) \in \mathcal{G}, \mathbf{v} = 0\} \\ &\geq \inf\{(\boldsymbol{\lambda}, 1)^T(\mathbf{v}, t) \mid (\mathbf{v}, t) \in \mathcal{G}\} \\ &= g(\boldsymbol{\lambda}), \end{aligned}$$

i.e., we have weak duality.

Now, we describe a variation on the geometric interpretation of duality in terms of \mathcal{G} , which explains why strong duality obtains for (most) convex problems. We define the set $\mathcal{A} \subseteq \mathbb{R}^n \times \mathbb{R}$ as

$$\mathcal{A} = \mathcal{G} + (\{0\} \times \mathbb{R}_+), \tag{6}$$

or, more explicitly,

$$\mathcal{A} = \{(\mathbf{v}, t) \mid \exists \mathbf{x} \in \text{dom} f, \mathbf{Ax} - \mathbf{b} = \mathbf{v}, f(\mathbf{x}) \leq t\},$$

We can think of \mathcal{A} as a sort of epigraph form of \mathcal{G} , since \mathcal{A} includes all the points in \mathcal{G} , as well as points that are ‘worse’, i.e., those with larger objective or inequality constraint function values. We can express the optimal value in terms of \mathcal{A} as

$$f^* = \inf\{t \mid (0, t) \in \mathcal{A}\}.$$

To evaluate the dual function at a point λ , we can minimize the affine function $(\lambda, 1)^T(\mathbf{v}, t)$ over \mathcal{A} , then

$$g(\lambda) = \inf\{(\lambda, 1)^T(\mathbf{v}, t) \mid (\mathbf{v}, t) \in \mathcal{A}\}.$$

If the infimum is finite, then

$$(\lambda, 1)^T(\mathbf{v}, t) \geq g(\lambda)$$

defines a nonvertical supporting hyperplane to \mathcal{A} . In particular, since $(0, f^*) \in \text{bd}(\mathcal{A})^3$, we have

$$f^* = (\lambda, 1)^T(0, f^*) \geq g(\lambda), \quad (7)$$

the weak duality lower bound. Strong duality holds if and only if we have equality in (7) for some dual feasible λ , i.e., there exists a nonvertical supporting hyperplane to \mathcal{A} at its boundary point $(0, f^*)$.

The set \mathcal{A} defined in (6) is readily shown to be convex if the underlying problem is convex. We define a second convex set \mathcal{B} as

$$\mathcal{B} = \{(0, s) \in \mathbf{R}^n \times \mathbf{R} \mid s < f^*\}.$$

The sets \mathcal{A} and \mathcal{B} do not intersect. To see this, suppose $(\mathbf{v}, t) \in \mathcal{A} \cap \mathcal{B}$. Since $(\mathbf{v}, t) \in \mathcal{B}$ we have $\mathbf{v} = 0$, and $t < f^*$. Since $(\mathbf{v}, t) \in \mathcal{A}$, there exists an \mathbf{x} with $\mathbf{A}\mathbf{x} - \mathbf{b} = 0$, and $f(\mathbf{x}) \leq t < f^*$, which is impossible since f^* is the optimal value of the primal problem. By the separating hyperplane theorem [3, Section 2.5.1], there exists $(\tilde{\lambda}, \mu) \neq 0$ and α such that

$$(\mathbf{v}, t) \in \mathcal{A} \implies \tilde{\lambda}^T \mathbf{v} + \mu t \geq \alpha, \quad (8)$$

$$(\mathbf{v}, t) \in \mathcal{B} \implies \tilde{\lambda}^T \mathbf{v} + \mu t \leq \alpha. \quad (9)$$

From (8) we conclude that $\mu \geq 0$. (Otherwise μt is unbounded below over \mathcal{A} , contradicting (8).) The condition (9) simply means that $\mu t \leq \alpha$ for all $t < f^*$, and hence, $\mu f^* \leq \alpha$. Together with (8) we conclude that for any $\mathbf{x} \in \text{dom}f$,

$$\tilde{\lambda}^T(\mathbf{A}\mathbf{x} - \mathbf{b}) + \mu f(\mathbf{x}) \geq \alpha \geq \mu f^*. \quad (10)$$

Assume that $\mu > 0$. In that case, since we can express the Lagrangian function as $L(\mathbf{x}, \lambda) = f(\mathbf{x}) + \lambda^T(\mathbf{A}\mathbf{x} - \mathbf{b})$, we can divide (10) by μ to obtain

$$L(\mathbf{x}, \tilde{\lambda}/\mu) \geq f^*$$

for all $\mathbf{x} \in \text{dom}f$, from which it follows, by minimizing over \mathbf{x} , that $g(\lambda) \geq f^*$, where we define

$$\lambda = \tilde{\lambda}/\mu.$$

By weak duality we have $g(\lambda) \leq f^*$, so in fact $g(\lambda) = f^*$. This shows that strong duality holds, and that the dual optimum is attained, at least in the case when $\mu > 0$. Now consider the case $\mu = 0$. From (10), we conclude that for all $\mathbf{x} \in \text{dom}f$,

$$\tilde{\lambda}^T(\mathbf{A}\mathbf{x} - \mathbf{b}) \geq 0. \quad (11)$$

From $(\tilde{\lambda}, \mu) \neq 0$ and $\mu = 0$, we conclude that $\tilde{\lambda} \neq 0$. Then (11) implies that for all $\mathbf{x} \in \text{dom}f$, $\tilde{\lambda}^T(\mathbf{A}\mathbf{x} - \mathbf{b}) \geq 0$. But $\tilde{\lambda}$ satisfies $\tilde{\lambda}^T(\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}) = 0$, and since $\tilde{\mathbf{x}} \in \text{int dom}f$, there are points in $\text{dom}f$ with $\tilde{\lambda}^T(\mathbf{A}\mathbf{x} - \mathbf{b}) < 0$ unless $\mathbf{A}^T \tilde{\lambda} = 0$. This, of course, contradicts our assumption that $\text{rank}(\mathbf{A}) = n$. \square

Remark. We have proven two crucial concepts from duality theory. To summarize them,

- By definition of f^* and d^* , we always have $d^* \leq f^*$ (**weak duality**).
- If a primal solution exists and Slater’s condition holds, we have $d^* = f^*$ (**strong duality**).

³ $\text{bd}(\mathcal{A})$ is the boundary of set \mathcal{A} .

Example 7. (Slater's condition) Let us consider solving $\min_{\mathbf{x} \in \mathcal{D}_\alpha} f(\mathbf{x})$ and so the feasible set is $\mathcal{D}_\alpha := \mathcal{X} \cap \mathcal{A}_\alpha$, where

$$\mathcal{X} := \{\mathbf{x} \in \mathbb{R}^2 : x_1^2 + x_2^2 \leq 1\}, \quad \mathcal{A}_\alpha := \{\mathbf{x} \in \mathbb{R}^2 : x_1 + x_2 = \alpha\},$$

where $\alpha \in \mathbb{R}$. Let's observe two cases where Slater's condition holds ($\alpha = 1/2$) and does not hold ($\alpha = \sqrt{2}$). We are looking for a feasible point that satisfies the (convex) inequality constraints *strictly*. First of all, \mathcal{X} is a convex set since $x_1^2 + x_2^2 - 1$ is a convex function of \mathbf{x} . \mathcal{A}_α contains only one affine constraint, hence it is also convex. When $\alpha = 1/2$, we can easily find a feasible point; e.g., $\mathbf{x} = (1/4, 1/4)$ that would satisfy the constraint in \mathcal{X} : $\frac{1}{4^2} + \frac{1}{4^2} < 1$ strictly. Hence, Slater's condition holds. On the other hand, when $\alpha = \sqrt{2}$, the only feasible point is $\mathbf{x} = (1/\sqrt{2}, 1/\sqrt{2})$, which yields $\frac{1}{(\sqrt{2})^2} + \frac{1}{(\sqrt{2})^2} = 1$. Therefore, it does not hold. More geometric intuition can be given in Figure 7, where it becomes evident that the figure on the right demonstrates an empty set for the relative interior.




Figure 6: $\mathcal{D}_{1/2}$ satisfies Slater's condition – $\mathcal{D}_{\sqrt{2}}$ does not satisfy Slater's condition.

5 Performance of optimization algorithms

In general, computing an **exact solution \mathbf{x}^*** is **impracticable**. Therefore, in the unconstrained setting, algorithms usually seek \mathbf{x}_ϵ^* that **approximates \mathbf{x}^*** up to $\epsilon > 0$ in some sense; i.e., finding \mathbf{x}_ϵ^* such that $f(\mathbf{x}_\epsilon^*) - f^* \leq \epsilon$. In constrained optimization problems, the pursuit of approximate solutions introduces a nuanced perspective that differs from the unconstrained case.

Consider a scenario where the objective is to minimize a simple convex function under certain constraints. In cases where you employ an algorithm – one that consistently operates within the feasible set – you can indeed attain ϵ -optimal solutions, similar to the unconstrained setting. This implies that the solutions generated by the algorithm will approach the optimal objective value within the ϵ boundary.

On the other hand, the situation is different when you have an algorithm that, alongside minimizing the objective function, also tries to maintain feasibility but does not rigidly enforce it at all times. In such instances, it is possible for a sequence of iterations to exhibit strange behavior: the objective values of iterates persistently may remain below the true minimum, even though our goal is to minimize the function. This phenomenon can be attributed to the algorithm's ability to process infeasible points, and then enforce feasibility which can initially yield a smaller objective value. This trajectory allows the algorithm to approach the minimum from below, showing that the path to approximate optimality in constraint problems can be far more complicated than in the unconstrained case because the difference $f(\mathbf{x}_\epsilon^*) - f^*$ can be negative. Therefore, we need an alternative performance metric for such algorithms. A possible metric might be the **time-to-reach ϵ** :

$$\text{time-to-reach } \epsilon = \text{number of iterations to reach } \epsilon \times \text{per iteration time}$$

Yet, still **the notion of ϵ -accuracy is elusive in constrained optimization!** We need a better characterization of ϵ -accurate solutions for the constrained optimization, which we cover in the next section.

5.1 Numerical ϵ -accuracy

The concept of ϵ -accuracy is vague in constrained setting. Therefore, we make a new definition of ϵ -accuracy.

Definition 5. (ϵ -accurate solutions [10]) Given a numerical *tolerance* $\epsilon \geq 0$, a point $\mathbf{x}_\epsilon^* \in \mathbb{R}^p$ is called an ϵ -*solution* of (3) if

$$\begin{cases} f(\mathbf{x}_\epsilon^*) - f^* & \leq \epsilon \text{ (objective residual),} \\ \|\mathbf{A}\mathbf{x}_\epsilon^* - \mathbf{b}\| & \leq \epsilon \text{ (feasibility gap),} \end{cases}$$

- When \mathbf{x}^* is unique, we can also obtain $\|\mathbf{x}_\epsilon^* - \mathbf{x}^*\| \leq \epsilon$ (iterate residual).

That is, we also add a ϵ tolerance to the feasibility gap⁴. Now, let's delve into another important aspect of assessing solution quality—the concept of “duality gap.”

5.2 Duality gap

To assess the quality of solutions, one approach is to measure the disparity in objective values between the primal and dual problems, known as *duality gap*. It is particularly interesting since it is 0 at the optimum for most convex problems. In a formal sense, we can consider the duality gap between a primal solution $\bar{\mathbf{x}}$ and a dual solution $\bar{\mathbf{y}}$ produced by our algorithm, as a ‘favorable’ value if it satisfies

$$\text{Gap}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = \max_{\mathbf{y} \in \mathcal{Y}} \Phi(\bar{\mathbf{x}}, \mathbf{y}) - \min_{\mathbf{x} \in \mathcal{X}} \Phi(\mathbf{x}, \bar{\mathbf{y}}) \leq \epsilon.$$

Now, we discuss an important characteristic of this gap. To this end, we begin by defining a concept known as the *saddle point*.

Definition 6. (*Saddle point*) We say $(\mathbf{x}^*, \mathbf{y}^*)$ is a *primal-dual solution* corresponding to primal and dual problems

$$f^* := \begin{cases} \min_{\mathbf{x} \in \mathbb{R}^p} & f(\mathbf{x}) \\ \text{s.t.} & \mathbf{A}\mathbf{x} = \mathbf{b}, \end{cases} \quad \text{and} \quad d^* := \max_{\mathbf{y} \in \mathbb{R}^n} d(\mathbf{y}) = \max_{\mathbf{y} \in \mathbb{R}^n} \min_{\mathbf{x}} \Phi(\mathbf{x}, \mathbf{y}).$$

if it is a *saddle point* of $\Phi(\mathbf{x}, \mathbf{y}) = f(\mathbf{x}) + \langle \mathbf{y}, \mathbf{A}\mathbf{x} - \mathbf{b} \rangle$:

$$\Phi(\mathbf{x}^*, \mathbf{y}) \leq \Phi(\mathbf{x}^*, \mathbf{y}^*) \leq \Phi(\mathbf{x}, \mathbf{y}^*), \quad \forall \mathbf{x} \in \mathbb{R}^p, \mathbf{y} \in \mathbb{R}^n.$$

Proposition 5.1. $\text{Gap}(\mathbf{x}, \mathbf{y}) \geq 0$ for any primal-dual feasible (\mathbf{x}, \mathbf{y}) and $\text{Gap}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = 0$ if and only if $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ is a *saddle point*.

Proof. By definition of a saddle point $(\mathbf{x}^*, \mathbf{y}^*)$, we have $\forall \mathbf{x} \in \mathbb{R}^p, \forall \mathbf{y} \in \mathbb{R}^n$:

$$\Phi(\mathbf{x}^*, \mathbf{y}) \stackrel{(*)}{\leq} \Phi(\mathbf{x}^*, \mathbf{y}^*) \stackrel{(**)}{\leq} \Phi(\mathbf{x}, \mathbf{y}^*).$$

Then, we can show nonnegativity of Gap:

$$\begin{aligned} \text{Gap}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) &= \max_{\mathbf{y} \in \mathcal{Y}} \Phi(\bar{\mathbf{x}}, \mathbf{y}) - \min_{\mathbf{x} \in \mathcal{X}} \Phi(\mathbf{x}, \bar{\mathbf{y}}) \\ &\geq \Phi(\bar{\mathbf{x}}, \mathbf{y}^*) - \min_{\mathbf{x} \in \mathcal{X}} \Phi(\mathbf{x}, \bar{\mathbf{y}}), \quad \text{by the definition of maximization} \\ &\geq \Phi(\mathbf{x}^*, \mathbf{y}^*) - \min_{\mathbf{x} \in \mathcal{X}} \Phi(\mathbf{x}, \bar{\mathbf{y}}), \quad \text{by the inequality (**)} \\ &\geq \Phi(\mathbf{x}^*, \bar{\mathbf{y}}) - \min_{\mathbf{x} \in \mathcal{X}} \Phi(\mathbf{x}, \bar{\mathbf{y}}), \quad \text{by the inequality (*)} \\ &\geq 0, \quad \text{by the definition of minimization.} \end{aligned}$$

If $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = (\mathbf{x}^*, \mathbf{y}^*)$, then all the inequalities will be equalities and $\text{Gap}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = 0$. □

The result of the proof demonstrates the connection between the duality gap and the existence of a saddle point and that the gap is always nonnegative. Specifically, when the duality gap is equal to zero, it implies the existence of a saddle point in the optimization problem, indicating a state where both primal and dual problems are optimally solved and their objective values coincide.

Now, let's move on to explore another crucial aspect of optimization, known as the Karush-Kuhn-Tucker (KKT) conditions.

⁴Note that the tolerances for objective, feasibility gap, or the iterate residual can be different. For the sake of simplicity, we use the same ϵ here.

5.3 Karush-Kuhn-Tucker (KKT) conditions

Under our assumptions, an equivalent characterization of $(\mathbf{x}^*, \mathbf{y}^*)$ is via the KKT conditions of the problem

$$\min_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b},$$

which reads

$$\begin{cases} 0 \in \partial_{\mathbf{x}} \Phi(\mathbf{x}^*, \mathbf{y}^*) = \mathbf{A}^T \mathbf{y}^* + \partial f(\mathbf{x}^*), \\ 0 = \nabla_{\mathbf{y}} \Phi(\mathbf{x}^*, \mathbf{y}^*) = \mathbf{Ax}^* - \mathbf{b}. \end{cases}$$

Namely, the subgradient of the Lagrangian function $L(\mathbf{x}, \mathbf{y}) = f(\mathbf{x}) + \mathbf{y}^T(\mathbf{Ax} - \mathbf{b})$ with respect to primal and dual variables should be in the subdifferential at optimality. With respect to \mathbf{x} , we get $0 \in \partial_{\mathbf{x}} \Phi(\mathbf{x}^*, \mathbf{y}^*) = \mathbf{A}^T \mathbf{y}^* + \partial f(\mathbf{x}^*)$, since f is nonsmooth. We have an affine function with respect to \mathbf{y} , hence it is smooth and we have $0 = \nabla_{\mathbf{y}} \Phi(\mathbf{x}^*, \mathbf{y}^*) = \mathbf{Ax}^* - \mathbf{b}$.

Duality and KKT conditions are fundamental concepts when it comes to designing an algorithm and evaluating its performance in the context of constrained optimization problems. In the following section, we will present a basic algorithm for tackling such problems.

6 Primal approach: The penalty method

One simple way to deal with constrained optimization is to solve its regularized version. That is, we convert constrained problem (**difficult**) to unconstrained (**easy**) one by penalizing feasibility with penalty parameter $\mu > 0$:

$$F_{\mu}(\mathbf{x}) := \left\{ f(\mathbf{x}) + \frac{\mu}{2} \|\mathbf{Ax} - \mathbf{b}\|^2 \right\} \xLeftrightarrow{\mu \rightarrow \infty} \min_{\mathbf{x} \in \mathbb{R}^p} \left\{ f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b} \right\}.$$

By selecting an appropriate penalty parameter, one that makes $\mathbf{Ax} - \mathbf{b}$ approach zero at optimality, we can achieve the same solution as the constrained approach. If $\mathbf{Ax} - \mathbf{b}$ is not exactly zero but is relatively small, we can still obtain an approximate solution. The value of μ essentially acts as a weighting factor that allows us to balance the minimization of the objectives, $f(\mathbf{x})$ and $\|\mathbf{Ax} - \mathbf{b}\|^2$, as per our requirements. As μ approaches 0, the problem focuses exclusively on minimizing $f(\mathbf{x})$, whereas as μ tends to infinity, it enforces $\mathbf{Ax} - \mathbf{b}$ to become exactly 0, reducing to the original constrained problem. This trade-off is visualized in Figure 7.




Figure 7: Trade-off between $\|\mathbf{Ax} - \mathbf{b}\|$ and $f(\mathbf{x})$ depending on the value of penalty μ .

Furthermore, we can explore alternative functions other than the **quadratic** $\frac{1}{2} \|\cdot\|^2$ e.g., exact nonsmooth penalty functions: $\mu \|\mathbf{Ax} - \mathbf{b}\|_2$ or $\mu \|\mathbf{Ax} - \mathbf{b}\|_1$. While these options have been proven to yield the exact solution when μ is finite, they are difficult to solve [6, Section 17.2], [1]. The quadratic one we use here is called the Quadratic Penalty (QP) method.

Notice that $F_{\mu}(\mathbf{x})$ is in the composite form, thus we can use algorithms that we have covered for the unconstrained nonsmooth case such as the proximal method, acceleration, etc. Intuitively, one can follow a very naive approach to find an approximate solution for the constrained problem as follows. We pick an initial penalty parameter μ_0 and solve the problem with one of these methods (e.g. proximal). After obtaining a solution, the penalty parameter is gradually increased, which effectively increases the penalty for constraint violations. Despite its simplicity, this method can be

useful in some cases and it lays the foundation for more advanced algorithms (as discussed in the next section), making it a valuable tool in optimization. The pseudo-code for this can be found below:

Quadratic penalty method (QP):
<ol style="list-style-type: none"> 1. Choose $\mathbf{x}_0 \in \mathbb{R}^p$ and $\mu_0 > 0$. 2. For $k = 0, 1, \dots$, perform: <ol style="list-style-type: none"> 2.a. $\mathbf{x}_k := \arg \min_{\mathbf{x} \in \mathbb{R}^p} \left\{ f(\mathbf{x}) + \frac{\mu_k}{2} \ \mathbf{Ax} - \mathbf{b}\ ^2 \right\}$. 2.b. Update $\mu_{k+1} > \mu_k$.

In the following theorem, we give a convergence property of the QP method.

Theorem 6.1. [6, Theorem 17.1] *Assume that f is smooth and $\mu_k \rightarrow \infty$. Then, every limit point $\bar{\mathbf{x}}$ of the sequence $\{\mathbf{x}_k\}$ is a solution of the constrained problem*

$$\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{R}^p} \{f(\mathbf{x}) : \mathbf{Ax} = \mathbf{b}\}.$$

Proof. Let \mathbf{x}^* be a global solution of the above problem, that is,

$$f(\mathbf{x}^*) \leq f(\mathbf{x}) \quad \text{for all } \mathbf{x} \text{ with } \mathbf{Ax} = \mathbf{b}.$$

Since \mathbf{x}_k minimizes $Q(\mathbf{x}; \mu_k) := f(\mathbf{x}) + \frac{\mu_k}{2} \|\mathbf{Ax} - \mathbf{b}\|^2$ for each k , we have that $Q(\mathbf{x}_k; \mu_k) \leq Q(\mathbf{x}^*; \mu_k)$, which leads to the inequality

$$f(\mathbf{x}_k) + \frac{\mu_k}{2} \|\mathbf{Ax}_k - \mathbf{b}\|^2 \leq f(\mathbf{x}^*) + \frac{\mu_k}{2} \|\mathbf{Ax}^* - \mathbf{b}\|^2 = f(\mathbf{x}^*). \quad (12)$$

By rearranging this expression, we obtain

$$\|\mathbf{Ax}_k - \mathbf{b}\|^2 \leq \frac{2}{\mu_k} [f(\mathbf{x}^*) - f(\mathbf{x}_k)]. \quad (13)$$

Suppose that $\bar{\mathbf{x}}$ is a limit point of $\{\mathbf{x}_k\}$, so that there is an infinite subsequence \mathcal{K} such that

$$\lim_{k \in \mathcal{K}} \mathbf{x}_k = \bar{\mathbf{x}}.$$

By taking the limit as $k \rightarrow \infty, k \in \mathcal{K}$, on both sides of (13), we obtain

$$\|\mathbf{A}\bar{\mathbf{x}} - \mathbf{b}\|^2 = \lim_{k \in \mathcal{K}} \|\mathbf{Ax}_k - \mathbf{b}\|^2 \leq \lim_{k \in \mathcal{K}} \frac{2}{\mu_k} [f(\mathbf{x}^*) - f(\mathbf{x}_k)] = 0,$$

where the last equality follows from $\mu_k \uparrow \infty$. Therefore, we have that $\mathbf{A}\bar{\mathbf{x}} - \mathbf{b} = 0$, so that $\bar{\mathbf{x}}$ is feasible. Moreover, by taking the limit as $k \rightarrow \infty$ for $k \in \mathcal{K}$ in (12), we have by nonnegativity of μ_k and of $\|\mathbf{Ax}_k - \mathbf{b}\|^2$ that

$$f(\bar{\mathbf{x}}) \leq f(\bar{\mathbf{x}}) + \lim_{k \in \mathcal{K}} \frac{\mu_k}{2} \|\mathbf{Ax}_k - \mathbf{b}\|^2 \leq f(\mathbf{x}^*)$$

Since $\bar{\mathbf{x}}$ is a feasible point whose objective value is no larger than that of the global solution \mathbf{x}^* , we conclude that $\bar{\mathbf{x}}$, too, is a global solution, as claimed. \square

This result requires us to find the global minimizer for each subproblem k and the minimization problems of step 2.a. of the algorithm become ill-conditioned as $\mu_k \rightarrow \infty$. Therefore, achieving this convergence to the global solution for constrained problems is typically not possible in practice. The ways to overcome these limitations could be to

- solve the subproblem inexactly, i.e., up to ϵ accuracy, or
- **linearize** it to simplify subproblems.

We, now, analyze the latter.

Generalized quadratic penalty method:	
1.	Choose $\mathbf{x}_0 \in \mathbb{R}^p$, $\mu_0 > 0$ and positive semidefinite matrix \mathbf{Q}_k .
2.	For $k = 0, 1, \dots$, perform:
2.a.	$\mathbf{x}_k := \arg \min_{\mathbf{x} \in \mathbb{R}^p} \left\{ f(\mathbf{x}) + \frac{\mu_k}{2} \ \mathbf{A}\mathbf{x} - \mathbf{b}\ ^2 + \frac{1}{2} \ \mathbf{x} - \mathbf{x}_{k-1}\ _{\mathbf{Q}_k}^2 \right\}$.
2.b.	Update $\mu_{k+1} > \mu_k$.

6.1 Quadratic penalty: Linearization

In the above pseudo-code, we see a challenging subproblem in step 2.a. To simplify this, we minimize a majorizer, which provides an upper bound on the original objective function. This majorizer is parameterized by the matrix \mathbf{Q}_k . Moreover, the choice of \mathbf{Q}_k and the way to update the penalty can significantly impact the algorithm's behavior. To address these, let us examine them more closely.

- **Minimize a majorizer:** Let $F_{\mu_k}(\mathbf{x}) := f(\mathbf{x}) + \frac{\mu_k}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$, then in step 2.a., we will minimize a *majorizer* of $F_{\mu_k}(\mathbf{x})$. That is, recall the inequality for L -smooth functions h : $h(\mathbf{x}) \leq h(\mathbf{x}_k) + \nabla h(\mathbf{x}_k)^\top (\mathbf{x} - \mathbf{x}_k) + \frac{L}{2} \|\mathbf{x} - \mathbf{x}_k\|^2$, which we refer to as quadratic majorizer for h (see Lecture 7 for details). Also, notice that $\frac{\mu_k}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$ is a L -smooth with $L = \mu_k \|\mathbf{A}\|^2$. Then, we can define our majorizer

$$\begin{aligned} F_{\mu_k}(\mathbf{x}) &= f(\mathbf{x}) + \frac{\mu_k}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \\ &\leq f(\mathbf{x}) + \frac{\mu_k}{2} \|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|^2 + \mu_k \langle \mathbf{A}^\top (\mathbf{A}\mathbf{x}_k - \mathbf{b}), \mathbf{x} - \mathbf{x}_k \rangle + \frac{\mu_k \|\mathbf{A}\|^2}{2} \|\mathbf{x} - \mathbf{x}_k\|^2 \\ &=: F_{\mu_k}^{\mathbf{x}_k}(\mathbf{x}). \end{aligned}$$

Then, we minimize the upper bound $F_{\mu_k}^{\mathbf{x}_k}(\mathbf{x})$ instead of $F_{\mu_k}(\mathbf{x})$.

$$\begin{aligned} \mathbf{x}_{k+1} &= \arg \min_{\mathbf{x}} F_{\mu_k}^{\mathbf{x}_k}(\mathbf{x}) \\ &= \text{prox}_{\frac{1}{\mu_k \|\mathbf{A}\|^2} f} \left(\mathbf{x}_k - \frac{1}{\|\mathbf{A}\|^2} \mathbf{A}^\top (\|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|) \right). \end{aligned}$$

This yields only one proximal operator instead of solving a difficult subproblem! Now, thanks to the addition of the semi-norm term, we can also control the step size by carefully selecting \mathbf{Q}_k and *accelerate* the algorithm.

- **Matrix \mathbf{Q}_k Choice:** The choice of the matrix \mathbf{Q}_k is critical. Different choices lead to different behaviors of the optimization method. For example:
 - $\mathbf{Q}_k = \mathbf{0}$ gives the standard QP.
 - $\mathbf{Q}_k = \mathbf{I}$ gives strongly convex subproblems.
 - $\mathbf{Q}_k = \alpha_k \mathbf{I} - \mu_k \mathbf{A}^\top \mathbf{A}$, with $\alpha_k \geq \mu_k \|\mathbf{A}\|^2$ gives

$$\mathbf{x}_k = \text{prox}_{\frac{1}{\alpha_k} f} \left(\mathbf{x}_{k-1} - \frac{\mu_k}{\alpha_k} \mathbf{A}^\top (\mathbf{A}\mathbf{x}_{k-1} - \mathbf{b}) \right),$$

and picking $\alpha_k = \mu_k \|\mathbf{A}\|^2$ gives

$$\mathbf{x}_k = \text{prox}_{\frac{1}{\mu_k \|\mathbf{A}\|^2} f} \left(\mathbf{x}_{k-1} - \frac{1}{\|\mathbf{A}\|^2} \mathbf{A}^\top (\mathbf{A}\mathbf{x}_{k-1} - \mathbf{b}) \right).$$

Notice that with careful choice of \mathbf{Q}_k and α_k , we just need to apply **only one proximal operator**, making it amenable to practical implementation. The ease of using the proximal operator in this context can be attributed to its fundamental properties, which are as follows:

- *single valued & non-expansive* since f is a proper convex function.

Linearized QP method (LQP)	Accelerated linearized QP method (ALQP)
<p>1. Choose $\mathbf{x}_0 \in \mathbb{R}^p$, $\sigma_0 = 1$, $\mu_0 > 0$.</p> <p>2. For $k = 0, 1, \dots$:</p> <p>2.a. $\mathbf{x}_{k+1} := \text{prox}_{\frac{1}{\mu_k \ \mathbf{A}\ ^2} f} \left(\mathbf{x}_k - \frac{1}{\ \mathbf{A}\ ^2} \mathbf{A}^\top (\mathbf{A} \mathbf{x}_k - \mathbf{b}) \right)$.</p> <p>2.b. Update σ_{k+1} s.t. $\frac{(1-\sigma_{k+1})^2}{\sigma_{k+1}} = \frac{1}{\sigma_k}$.</p> <p>2.c. Update $\mu_{k+1} = \sqrt{\sigma_{k+1}}$.</p>	<p>1. Choose $\mathbf{x}_0, \mathbf{y}_0 \in \mathbb{R}^p$, $\tau_0 = 1$, $\mu_0 > 0$.</p> <p>2. For $k = 0, 1, \dots$:</p> <p>2.a. $\mathbf{x}_{k+1} := \text{prox}_{\frac{1}{\mu_k \ \mathbf{A}\ ^2} f} \left(\mathbf{y}_k - \frac{1}{\ \mathbf{A}\ ^2} \mathbf{A}^\top (\mathbf{A} \mathbf{y}_k - \mathbf{b}) \right)$.</p> <p>2.b. $\mathbf{y}_{k+1} := \mathbf{x}_{k+1} + \frac{\tau_{k+1}(1-\tau_k)}{\tau_k} (\mathbf{x}_{k+1} - \mathbf{x}_k)$.</p> <p>2.c. Update $\mu_{k+1} = \mu_k(1 + \tau_{k+1})$.</p> <p>2.d. Update $\tau_{k+1} \in (0, 1)$ as the unique positive root of $\tau^3 + \tau^2 + \tau_k^2 \tau - \tau_k^2 = 0$.</p>

Table 1: Comparison of LQP and ALQP methods.

- *distributes* when the primal problem has *decomposable* structure:

$$f(\mathbf{x}) := \sum_{i=1}^m f_i(\mathbf{x}_i), \quad \text{and} \quad \mathcal{X} := \mathcal{X}_1 \times \dots \times \mathcal{X}_m.$$

where $m \geq 1$ is the number of components.

- *often efficient & has closed form expression*. For instance, if $f(\mathbf{z}) = \|\mathbf{z}\|_1$, then the prox-operator performs coordinate-wise soft-thresholding by 1.

When we update the penalty parameter with Linear QP method (LQP) as described in Table 1(left), we achieve a \sqrt{k} convergence rate. Alternatively, we can employ an accelerated LQP (ALQP) as in Table 1(right) to convert this into a $\frac{1}{k}$ rate. The following theorem summarizes these convergence rates.

Theorem 6.2 (Convergence [11]). *The convergence rates for the LQP and ALQP methods are as follows:*

- *LQP:*

$$\begin{aligned} |f(\mathbf{x}_k) - f(\mathbf{x}^*)| &\leq O(\mu_0 k^{-1/2} + \mu_0^{-1} k^{-1/2}) \\ \|\mathbf{A} \mathbf{x}_k - \mathbf{b}\| &\leq O(\mu_0^{-1} k^{-1/2}) \end{aligned}$$

- *ALQP:*

$$\begin{aligned} |f(\mathbf{x}_k) - f(\mathbf{x}^*)| &\leq O(\mu_0 k^{-1} + \mu_0^{-1} k^{-1}) \\ \|\mathbf{A} \mathbf{x}_k - \mathbf{b}\| &\leq O(\mu_0^{-1} k^{-1}) \end{aligned}$$

6.1.1 LQP and ALQP in practice

In practice, these methods rarely work better than the worst case. To illustrate this, consider a nonsmooth problem: Square-root Least Absolute Shrinkage and Selection Operator (SQRT Lasso), defined as

$$\min_{\mathbf{x} \in \mathbb{R}^p} \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2 + \lambda \|\mathbf{x}\|_1.$$

Figure 8 presents the performance of ALQP on this problem. By visual inspection, we can easily see the $\frac{1}{k}$ convergence rate (starred blue line). However, its performance is still very close to the worst case scenarios. Thus, we will explore algorithms with superior performance in Lecture 14.

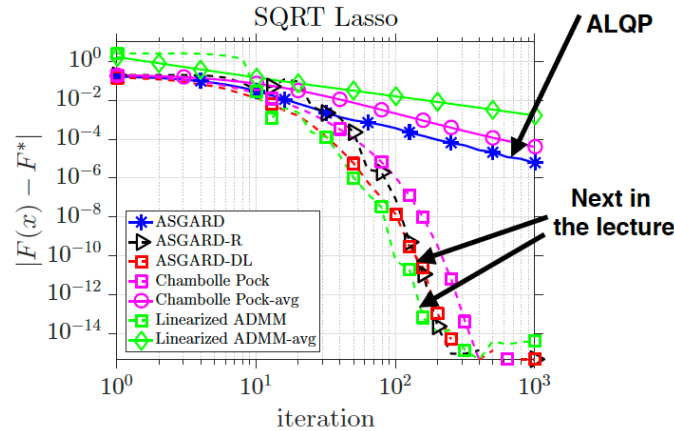


Figure 8: The performance of ALQP on SQR T Lasso.

References

- [1] Dimitri P Bertsekas. Necessary and sufficient conditions for a penalty method to be exact. *Mathematical Programming*, 9(1):87–99, 1975.
- [2] Jonathan Borwein and Adrian Lewis. *Convex Analysis and Nonlinear Optimization*. Springer, 2006.
- [3] Stephen P Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [4] Sham M Kakade, Shai Shalev-Shwartz, and Ambuj Tewari. Regularization techniques for learning with matrices. *Journal of Machine Learning Research*, 13(59):1865–1890, 2012.
- [5] GM Korpelevich. An extragradient method for finding saddle points and for other problems. *Ekonomika i Matematicheskie Metody*, 12(4):747–756, 1976.
- [6] Jorge Nocedal and Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, 2 edition, 1999.
- [7] R Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [8] Shai Shalev-Shwartz. *Online learning: Theory, algorithms, and applications*. PhD thesis, Hebrew University, 2007.
- [9] Shai Shalev-Shwartz and Yoram Singer. On the equivalence of weak learnability and linear separability: New relaxations and efficient boosting algorithms. *Machine Learning*, 80:141–163, 2010.
- [10] Quoc Tran-Dinh and Volkan Cevher. Constrained convex minimization via model-based excessive gap. *Advances in Neural Information Processing Systems*, 27, 2014.
- [11] Quoc Tran-Dinh, Olivier Fercoq, and Volkan Cevher. A smooth primal-dual optimization framework for nonsmooth composite convex minimization. *SIAM Journal on Optimization*, 28(1):96–134, 2018.