Chapter 5

Optimization

The foundation of engineering is the ability to use math and physics to design and optimize complex systems. The advent of computers has made this possible on an unprecedented scale. This chapter provides a brief introduction to optimization theory.

5.1 Derivatives in Banach Spaces

Definition 5.1.1. *Let $f\colon X \to Y$ be a mapping from a vector space X to a Banach space $(Y, \|\cdot\|)$. Then, if it exists, the* ***Gâteuax differential*** *of f at $\underline{x}$ in direction $\underline{h}$ is given by*

$$\delta f(\underline{x};\underline{h}) \triangleq \lim_{t\to 0} \frac{f(\underline{x}+t\underline{h}) - f(\underline{x})}{t},$$

where the limit is with respect to the implied mapping from $\mathbb{R}$ to Y.

Lemma 5.1.2. *Let $Y = (\mathbb{R}, |\cdot|)$ and suppose that $\delta f(\underline{x};\underline{h}) < 0$ exists for some f, $\underline{x}$, and $\underline{h}$. Then, there exists $t_0 > 0$ such that, for all $t \in (0, t_0)$, one has*

$$f(\underline{x}+t\underline{h}) < f(\underline{x}).$$

Proof. This follows from applying the definition of the limit. □

Problem 5.1.3. *Suppose $X = Y = L^1([0,1])$ is the Banach space of Lebesgue absolutely integrable functions mapping $[0,1]$ to $\mathbb{R}$ and $f(\underline{x}) = \|\underline{x}\| = \int_0^1 |x(s)|\mathrm{d}s$ is the norm of $\underline{x}$. Compute $\delta f(\underline{x};\underline{h})$ for this case by interchaging the limit and integration in*

$$\lim_{t\to 0} \int_0^1 \frac{1}{t}\left(|x(s)+th(s)| - |x(s)|\right)\mathrm{d}s.$$

Definition 5.1.4. *Let $f\colon X \to Y$ be a mapping from a vector space X to a Banach space $(Y, \|\cdot\|)$. Then, f is **Gâteuax differentiable** at $\underline{x}$ if the Gâteuax differential $\delta f(\underline{x}; \underline{h})$ exists for all $\underline{h} \in X$ and is a continuous linear function of $\underline{h}$.*

Definition 5.1.5. *Let $f\colon X \to Y$ be a mapping from a Banach space $(X, \|\cdot\|_X)$ to a Banach space $(Y, \|\cdot\|_Y)$. Then, f is **Fréchet differentiable** at $\underline{x}_0$ if there is a continuous linear transformation $T\colon X \to Y$ satisfying*

$$\lim_{\underline{h}\to\underline{0}} \frac{\|f(\underline{x}+\underline{h}) - f(\underline{x}) - T(\underline{h})\|_Y}{\|\underline{h}\|_X} = 0,$$

where the limit is with respect to the implied Banach space mapping $X \to \mathbb{R}$.

Example 5.1.6. *A function $f\colon \mathbb{R}^n \to \mathbb{R}^m$ with $f = (f_1, f_2, \ldots, f_m)^T$ is (Fréchet) differentiable at $\underline{x}_0$ if the mapping J from $\mathbb{R}^n$ to the **Jacobian matrix**,*

$$J(\underline{x}) = f'(\underline{x}) \triangleq \begin{bmatrix} \frac{\partial f_1}{\partial x_1}(\underline{x}) & \frac{\partial f_1}{\partial x_2}(\underline{x}) & \cdots & \frac{\partial f_1}{\partial x_n}(\underline{x}) \\ \frac{\partial f_2}{\partial x_1}(\underline{x}) & \frac{\partial f_2}{\partial x_2}(\underline{x}) & \cdots & \frac{\partial f_2}{\partial x_n}(\underline{x}) \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1}(\underline{x}) & \frac{\partial f_m}{\partial x_2}(\underline{x}) & \cdots & \frac{\partial f_m}{\partial x_n}(\underline{x}) \end{bmatrix},$$

exists and is continuous in $\underline{x}$ at $\underline{x} = \underline{x}_0$. A necessary and sufficient condition for this is that each partial derivative is continuous in $\underline{x}$ at $\underline{x} = \underline{x}_0$.

*If $m = 1$, then the Jacobian is also called the **gradient** of the function*

$$f'(\underline{x}) = \nabla f(\underline{x}) \triangleq \begin{bmatrix} \frac{\partial f}{\partial x_1}(\underline{x}) & \frac{\partial f}{\partial x_2}(\underline{x}) & \cdots & \frac{\partial f}{\partial x_n}(\underline{x}) \end{bmatrix}.$$

5.2 Unconstrained Optimization

Functions mapping elements of a vector space (over F) down to the scalar field F play a very special role in the analysis of vector spaces.

Definition 5.2.1. *Let V be a vector space over F. Then, a **functional** on V is a function $f\colon V \to F$ that maps V to F.*

Linear functionals (i.e., functionals that are linear) are used to define many important concepts in abstract vector spaces. For unconstrained optimization, however, linear functionals are not interesting because they are either zero or they achieve all values in F.

Definition 5.2.2. *Let $(X, \|\cdot\|)$ be a normed vector space. Then, a real functional $f\colon X \to \mathbb{R}$ achieves a **local minimum value** at $\underline{x}_0 \in X$ if there is an $\epsilon > 0$ such that, for all $\underline{x} \in X$ satisfying $\|\underline{x} - \underline{x}_0\| < \epsilon$, we have $f(\underline{x}) \geq f(\underline{x}_0)$. If the bound holds for all $x \in X$, then the local minimum is also a **global minimum value**.*

Theorem 5.2.3. *Let $(X, \|\cdot\|)$ be a normed vector space and $f\colon X \to \mathbb{R}$ be a real functional. If $\delta f(\underline{x}_0, \underline{h})$ exists and is negative for any $\underline{h} \in X$, then $\underline{x}_0$ is not a local minimum value.*

Proof. First, we apply Lemma 5.1.2 with the $\underline{x}$ and $\underline{h}$ for which $\delta f(\underline{x}_0, \underline{h}) < 0$. This gives a $t_0 > 0$ such that $f(\underline{x}_0 + t\underline{h}) < f(\underline{x}_0)$ for all $t \in (0, t_0)$. Thus, there can no be no $\epsilon > 0$ satisfying the definition of a local minimum value in Definition 5.2.2. □

5.3 Convex Functionals

Convexity is a particularly nice property of spaces and functionals that leads to well-defined minimum values.

Definition 5.3.1. *A Banach space X is called **strictly convex** if the unit ball, given by $\{x \in X \mid \|x\| \leq 1\}$, is a strictly convex set. An equivalent condition is that equality in the triangle inequality (i.e., $\|\underline{x} + \underline{y}\| = \|\underline{x}\| + \|\underline{y}\|$) for non-zero vectors implies that $\underline{x} = s\underline{y}$ for some $s \in F$.*

Definition 5.3.2. *Let V be a vector space, $A \subseteq V$ be a convex set, and $f\colon V \to \mathbb{R}$ be a functional. Then, the functional f is called **convex** on A if, for all $\underline{a}_1, \underline{a}_2 \in A$ and $\lambda \in (0, 1)$, we have*

$$f(\lambda \underline{a}_1 + (1 - \lambda)\underline{a}_2) \leq \lambda f(\underline{a}_1) + (1 - \lambda) f(\underline{a}_2).$$

*The functional is **strictly convex** if equality occurs only when $\underline{a}_1 = \underline{a}_2$.*

Example 5.3.3. *Let $(X, \|\cdot\|)$ be a normed vector space. Then, the norm $\|\cdot\|\colon X \to \mathbb{R}$ is a convex functional on X. Try proving this as an exercise.*

Theorem 5.3.4. *Let $(X, \|\cdot\|)$ be a normed vector space, $A \subseteq X$ be a convex set, and $f\colon X \to \mathbb{R}$ be a convex functional on A. Then, any local minimum value of f on A is a global minimum value on A. If the functional is strictly convex on A and achieves a local minimum value on A, then there is a unique point $\underline{x}_0 \in A$ that achieves the global minimum value on A.*

Proof. Let $\underline{x}_0 \in A$ a point where the functional achieves a local minimum value. Proving by contradiction, we suppose that there is another point $\underline{x}_1 \in A$ such that $f(\underline{x}_1) < f(\underline{x}_0)$. From the definition of a local minimum value, we find an $\epsilon > 0$ such that $f(\underline{x}) \geq f(\underline{x}_0)$ for all $\underline{x} \in A$ satisfying $\|\underline{x} - \underline{x}_0\| < \epsilon$. Choosing $\lambda < \frac{\epsilon}{\|\underline{x}_0 - \underline{x}_1\|}$ in $(0, 1)$ and $\underline{x} = (1-\lambda)\underline{x}_0 + \lambda \underline{x}_1$ implies that $\|\underline{x} - \underline{x}_0\| < \epsilon$ while the convexity of f implies that

$$f(\underline{x}) = f\left((1-\lambda)\underline{x}_0 + \lambda\underline{x}_1\right) \leq (1-\lambda)f(\underline{x}_0) + \lambda f(\underline{x}_1) < f(\underline{x}_0).$$

This contradicts the definition of a local minimum value and implies that $f(\underline{x}_0)$ is a global minimum value on A. If f is strictly convex and $f(\underline{x}_1) = f(\underline{x}_0)$, then we suppose that $\underline{x}_0 \neq \underline{x}_1$. In this case, strict convexity implies that

$$f\left((1-\lambda)\underline{x}_0 + \lambda\underline{x}_1\right) < (1-\lambda)f(\underline{x}_0) + \lambda f(\underline{x}_1) = f(\underline{x}_0).$$

This contradicts the fact that $f(\underline{x}_0)$ is a global minimum value on A and implies that $\underline{x}_0 = \underline{x}_1$ is unique. □

Theorem 5.3.5. *Let $(X, \|\cdot\|)$ be a normed vector space and $f\colon X \to \mathbb{R}$ be a convex functional on a convex set $A \subseteq X$. If f is Gâteuax differentiable at $\underline{x}_0 \in A$, then*

$$f(\underline{x}) \geq f(\underline{x}_0) + \delta f(\underline{x}_0; \underline{x} - \underline{x}_0)$$

for all $\underline{x} \in A$. If f is strictly convex then the inequality is strict for $\underline{x} \neq \underline{x}_0$.

Proof. By the convexity of A and f, we have $\underline{x}_0 + \lambda(\underline{x} - \underline{x}_0) \in A$ and

$$f\left(\underline{x}_0 + \lambda(\underline{x} - \underline{x}_0)\right) \leq f(\underline{x}_0) + \lambda\left(f(\underline{x}) - f(\underline{x}_0)\right) \tag{5.1}$$

for all $\lambda \in (0, 1)$. Also, if f is strictly convex, then (5.1) strict for $\underline{x} \neq \underline{x}_0$. Thus,

$$f(\underline{x}) \geq f(\underline{x}_0) + \frac{f\left(\underline{x} + \lambda(\underline{x} - \underline{x}_0)\right) - f(\underline{x}_0)}{\lambda}$$

and taking the limit at $\lambda \downarrow 0$ completes the proof for a convex functional.

For the case where f is strictly convex, we first apply the convex result to see

$$f(\underline{x}_0 + \lambda(\underline{x} - \underline{x}_0)) \geq f(\underline{x}_0) + \delta f(\underline{x}_0; \lambda(\underline{x} - \underline{x}_0)) = f(\underline{x}_0) + \lambda \delta f(\underline{x}_0; \underline{x} - \underline{x}_0),$$

where the second step holds because $\delta f(\underline{x}; \underline{h})$ is linear in $\underline{h}$. This gives

$$\delta f(\underline{x}_0; \underline{x} - \underline{x}_0) \leq \frac{f(\underline{x}_0 + \lambda(\underline{x} - \underline{x}_0)) - f(\underline{x}_0)}{\lambda} < f(\underline{x}) - f(\underline{x}_0),$$

where the second inequality holds because (5.1) is a strict inequality for $\underline{x} \neq \underline{x}_0$. □

Corollary 5.3.6. *Let $(X, \|\cdot\|)$ be a normed vector space and $f\colon X \to \mathbb{R}$ be a convex functional on a convex set $A \subseteq X$. If f is Gâteuax differentiable at $\underline{x}_0 \in A$ and $\delta f(\underline{x}_0; \underline{x} - \underline{x}_0) = 0$ for all $\underline{x} \in A$, then*

$$f(\underline{x}_0) = \min_{\underline{x} \in A} f(\underline{x}).$$

If f is strictly convex, $\underline{x}_0$ is the unique minimizer over A.

5.4 Constrained Optimization

Lagrangian optimization is an indispensable tool in engineering and physics that allows one to solve constrained non-linear optimization problems. For convex problems, there are now efficient algorithms that can handle thousands of variables and constraints. In some cases, there are also analytical techniques that allow one to derive tight bounds on optimum value. These approaches have become so common that convex Lagrangian optimization problems are now taught as a fundamental part of the graduate engineering curriculum. For simplicity, we focus on the case where the domain $\mathcal{D}$ is a subset of the finite-dimensional real space $\mathbb{R}^n$.

Constrained non-linear optimization problems over $\mathbb{R}^n$ can be put into the following **standard form**. Let $f_i\colon \mathcal{D} \to \mathbb{R}$ and $h_j\colon \mathcal{D} \to \mathbb{R}$ be a real functionals on $\mathcal{D} \subseteq \mathbb{R}^n$ for $i = 1, 2, \ldots, m$ and $j = 1, 2, \ldots, p$. Then, the standard form is

$$\begin{aligned} \text{minimize} \quad & f_0(\underline{x}) \\ \text{subject to} \quad & f_i(\underline{x}) \le 0, \quad i = 1, 2, \ldots, m \\ & h_j(\underline{x}) = 0, \quad j = 1, 2, \ldots, p. \end{aligned}$$

The function f_0 is called the **objective function** while the functions $f_1, \ldots, f_m$ are called inequality constraints and the functions $h_1, \ldots, h_p$ are called equality constraints.

Definition 5.4.1. *A vector $\underline{x} \in \mathcal{D}$ is **feasible** if it satisfies the constraints. Let $\mathcal{F} = \{\underline{x} \,|\, f_i(\underline{x}) \le 0, i = 1, 2, \ldots, m\,, h_j(\underline{x}) = 0, j = 1, \ldots, p\}$ be the set of feasible vectors. Then, the problem is feasible if $\mathcal{F} \neq \emptyset$.*

Definition 5.4.2. *The **optimal value** is*

$$p^* = \inf\left\{f_0(\underline{x}) \,|\, \underline{x} \in \mathcal{F}\right\}.$$

By convention, p^* is allowed to take infinite values and $p^* = \infty$ if the problem is not feasible.

Evaluating the function at any feasible point gives the trivial upper bound

$$p^* \leq f_0(\underline{x}) \; \forall \underline{x} \in \mathcal{F}.$$

The optimization of linear function with arbitrary affine equality and inequality constraints is called a **linear program**. Any linear program can be reduced to the following standard form by adding additional variables.

Definition 5.4.3. *The standard form of a linear program (LP) is given by*

$$\begin{aligned} \text{minimize} \quad & \underline{c}^T \underline{x} \\ \text{subject to} \quad & A\underline{x} = \underline{b} \\ & \underline{x} \succeq \underline{0}. \end{aligned}$$

5.4.1 The Lagrangian

The Lagrangian is used to transform constrained optimization problems into unconstrained optimization problems. One can think of it as introducing a non-negative cost λ_i (resp. ν_j) associated with violating the i-th inequality (resp. j-th equality) constraint.

Definition 5.4.4. *The **Lagrangian** $L \colon \mathcal{D} \times \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}$ associated with optimization problem is*

$$L(\underline{x}, \underline{\lambda}, \underline{\nu}) = f_0(\underline{x}) + \sum_{i=1}^{m} \lambda_i f_i(\underline{x}) + \sum_{j=1}^{p} \nu_j h_j(\underline{x}),$$

*where λ_i is the **Lagrange multiplier** associated with the i-th inequality constraint and ν_j is the Lagrange multiplier associated with the j-th equality constraint.*

Definition 5.4.5. *A point $\underline{x}^*$ is called **locally optimal** if there is an $\epsilon_0 > 0$ such that, for all $\epsilon < \epsilon_0$, it holds that $f_0(\underline{x}) \geq f_0(\underline{x}^*)$ for all $\underline{x} \in \mathcal{F}$ satisfying $\|\underline{x} - \underline{x}^*\| < \epsilon$. The i-th inequality constraint is **active** at $\underline{x}^*$ if $f_i(\underline{x}^*) = 0$. Otherwise, it is inactive.*

Theorem 5.4.6 (Karush-Kuhn-Tucker)**.** *Assume the functions f_i and h_j are continuously differentiable and let $A = \{i \in [m] \,|\, f_i(\underline{x}^*) = 0\}$ be the set of active constraints at $\underline{x}^*$. Then, $\underline{x}^*$ is locally optimal only if $\underline{\lambda}^* \geq 0$ and $\underline{\nu}^*$ exist such that*

$$\nabla f_0(\underline{x}^*) + \sum_{i\in A} \lambda_i^* \nabla f_i(\underline{x}^*) + \sum_{j=1}^{p} \nu_j^* \nabla h_j(\underline{x}^*) = \underline{0} \tag{5.2}$$

This theorem provides a necessary condition for a point $\underline{x}^*$ to be locally optimal for a constrained optimization problem. Before considering its proof, it is useful to discuss the geometric picture upon which it is based. First, we note that the negative gradient $-\nabla f_0(\underline{x}^*)$ gives the direction of steepest descent for the objective function.

Now, consider what happens if we evaluate the function at $\underline{x}(t) = \underline{x}^* + t\underline{y}$ for some direction $\underline{y}$ and a sufficiently small $t > 0$. For any continuously differentiable function f, the definition of the derivative implies that

$$f(\underline{x}(t)) = f(\underline{x}^*) + t\underline{y}^H \nabla f(\underline{x}^*) + o(t),$$

where $o(t) \to 0$ as $t \to 0$. If the problem is unconstrained (e.g., $m = p = 0$), then $\nabla f_0(\underline{x}^*)$ must be $\underline{0}$. Otherwise, one is guaranteed to reduce the function by choosing $\underline{y} = -\nabla f_0(\underline{x}^*)$ (e.g., see Lemma 5.1.2). If there are constraints, however, then $\underline{x}(t)$ may be infeasible. For the j-th equality constraint, the definition of the derivative implies that $\underline{x}(t)$ will be infeasible if $|\underline{y}^H \nabla h_j(\underline{x}^*)| > 0$. Thus, we certainly need $\underline{y}^H \nabla h_j(\underline{x}^*) = 0$ for all j.

If the i-th inequality constraint is active (i.e., $f_i(\underline{x}^*) = 0$), then the definition of the derivative implies that $\underline{x}(t)$ will be infeasible if $\underline{y}^H \nabla f_i(\underline{x}^*) > 0$. Thus, we certainly need $\underline{y}^H \nabla f_i(\underline{x}^*) \leq 0$ for all $i \in A$. If the constraint is inactive (i.e., $f_i(\underline{x}^*) < 0$), then due to continuity it will remain satisfied for small perturbations of $\underline{x}^*$.

The geometric picture implied by Theorem 5.4.6 is that of a game where one would like to decrease the objective $f_0(\underline{x}^*)$ choosing $\underline{y}$ such that $\underline{y}^H \nabla f_0(\underline{x}^*) < 0$ but there are constraints on the set of allowable $\underline{y}$'s. Let $H = \mathrm{span}(\{\nabla h_j(\underline{x}^*)\})$ be the subspace of directions that violate the equality constraints at $\underline{x}^*$. Similarly, let the cone of directions that violate the active inequality constraints is given by

$$F = \left\{ \sum_{i\in A} \lambda_i \nabla f_i(\underline{x}^*) \,\middle|\, \lambda_i \geq 0, i \in A \right\}.$$

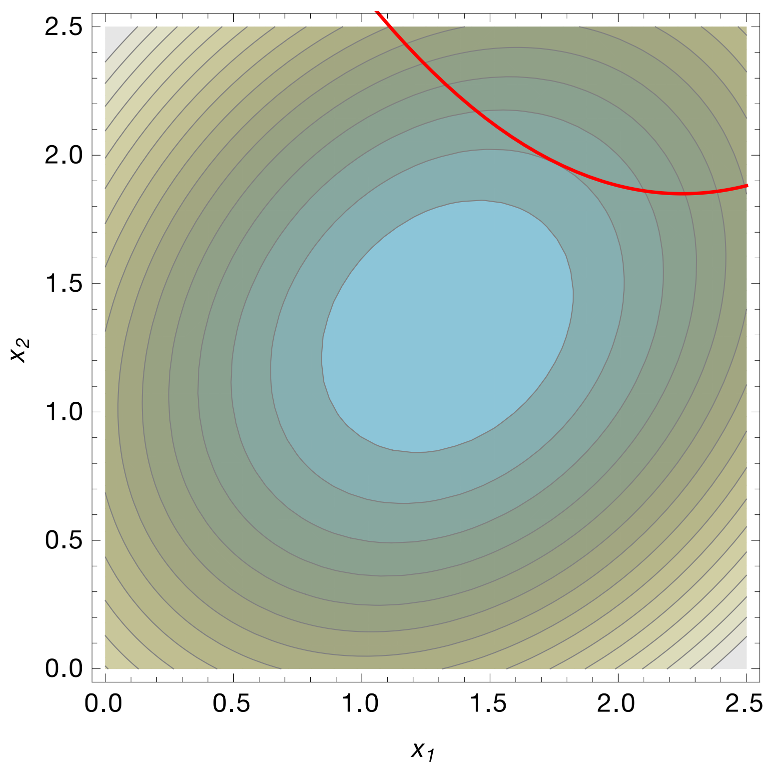

Figure 5.1: A contour plot of the function $f_0(x_1, x_2) = (x_1-1)^2+(x_2-1)^2-x_1x_2/2$ whose minimum occurs at $(4/3, 4/3)$ (i.e., the center of the blue ellipse). The red line indicates the inequality constraint $f_1(x_1, x_2) = 1.85+(x_1-2.25)^2/2-x_2 \leq 0$. The picture shows that the constrained minimum occurs at the intersection of the contour tangent line and the active constraint line.

Thus, one must at least pick a direction $\underline{y}$ that is orthogonal to all vectors in H and has a non-positive inner product with all vectors in F.

Let the matrix P define the orthogonal projection of $\mathbb{R}^n$ onto $H^\perp$. Using this, we can translate the equation (5.2) into the statement

$$-P\nabla f_0(\underline{x}^*) \in PF$$

or "the projection of the descent direction lies in the projection of the cone of directions that violate the inequality constraints". The reason for this is that we can absorb the ∇h_j terms into the ∇f_i terms by defining

$$\underline{f}^{(i)} = \nabla f_i(\underline{x}^*) + \sum_{j=1}^{p} \nu_{j,i} \nabla h_j(\underline{x}^*) = P\nabla f_i(\underline{x}^*)$$

so that $\underline{f}^{(i)} \in H^{\perp}$ for $i = 0, 1, \ldots, m$. Then, the cone PF is defined by

$$PF = \left\{ \sum_{i \in A} \lambda_i \underline{f}^{(i)} \,\middle|\, \lambda_i \geq 0, i \in A \right\}.$$

If $-P\nabla f_0(\underline{x}^*) \notin PF$, then we project $-P\nabla f_0(\underline{x}^*)$ onto PF to get a non-zero residual $\underline{y}$. The resulting vector gives a direction where the objective function decreases and the constraints remain *almost* satisfied. The challenge in making this proof precise is that, unless the equality constraints are affine, the constraints may not be exactly satisfied for $t > 0$. In standard profos of this result, this difficulty is overcome by using the implicit function theorem to construct an $\underline{x}(t)$ that starts in the direction on $\underline{y}$ but is perturbed slightly to remain feasible.

Proof. For simplicity, we prove only the case where $h_j(\underline{x}) = \underline{a}_j^H \underline{x} - \underline{b}$ is affine and PF does contain a line (i.e., $\{\alpha \underline{z} \,|\, \alpha \in \mathbb{R}\}$ for some $\underline{z}$). First, we define

$$\underline{y}(\underline{\lambda}, \underline{\nu}) = -\nabla f_0(\underline{x}^*) - \sum_{i=1}^{m} \lambda_i \nabla f_i(\underline{x}^*) - \sum_{j=1}^{p} \nu_j \underbrace{\nabla h_j(\underline{x}^*)}_{\underline{a}_j}.$$

The vector $\underline{y}(\underline{\lambda}, \underline{\nu})$ can be seen as the residual of the descent direction for the objective function after the constraint gradients have been used to cancel some parts. Next, we let $\underline{\nu}^*(\underline{\lambda}) = \arg\min_{\underline{\nu} \in \mathbb{R}^p} \|\underline{y}(\underline{\lambda}, \underline{\nu})\|_2$ and apply the best approximation theorem to see that

$$\underline{y}(\underline{\lambda}, \underline{\nu}^*(\underline{\lambda})) = P\underline{y}(\underline{\lambda}, \underline{\nu}),$$

where P is orthogonal projection onto $H^{\perp}$ and $H = \mathrm{span}(\{\underline{a}_j\})$. This ensures that each $h_j(\underline{x}^* + t\underline{y}(\underline{\lambda}, \underline{\nu}^*(\underline{\lambda}))) = 0$ for all $\underline{\lambda} \in \mathbb{R}^m$ and $t \in \mathbb{R}$.

Continuing, we define $\underline{y}^* = \arg\min_{\underline{\lambda} \in \mathbb{R}^m, \underline{\lambda} \geq \underline{0}} \|\underline{y}(\underline{\lambda}, \underline{\nu}^*(\underline{\lambda}))\|_2$. This implies that $\underline{y}^*$ is the error vector for the projection of $-P\nabla f_0(\underline{x}^*)$ onto the convex set PF. The projection itself is given by $\underline{z} = -P\nabla f_0 - \underline{y}^*$ and Lemma 4.6.4 shows that $\left(\underline{y}^*\right)^H (\underline{z}) \geq 0$. Using this, we see that

$$\left(\underline{y}^*\right)^H (-P\nabla f_0(\underline{x}^*)) = \left(\underline{y}^*\right)^H \left(\underline{z} + \underline{y}^*\right) \geq \|\underline{y}^*\|_2^2.$$

If (5.2) cannot be satisfied by some $\underline{\lambda} \geq \underline{0}$ and $\underline{\nu}$, then $\underline{y}^* \neq \underline{0}$ and $\|\underline{y}^*\|_2 > 0$. This shows that $\underline{y}^*$ points in a direction that decreases the value of the objective function.

But, the $\underline{y}^*$ direction is only guaranteed to preserve feasibility to first order (i.e., $(\underline{y}^*)^H P\nabla f_i(\underline{x}^*) \leq 0$). To fix this, one can add to $\underline{y}^*$ a sufficiently small vector $\underline{w}$

satisfying $\underline{w}^H P \nabla f_i(\underline{x}^*) < 0$ for all $i = 1, 2, \ldots, m$. Such a $\underline{w}$ lies in the "interior of the polar cone of PF" and will exist as long as PF does not contain a line. With this modification, the definition of the derivative implies that, for sufficiently small t, $\underline{x}(t) = \underline{x}^* + t(\underline{y}^* + \underline{w})$ will be a feasible vector satisfying $f_0(\underline{x}(t)) < f_0(\underline{x}^*)$. □

5.4.2 Lagrangian Duality

Definition 5.4.7. *The* ***Lagrangian dual*** *function is defined to be*

$$g(\underline{\lambda}, \underline{\nu}) = \inf_{\underline{x} \in \mathcal{D}} L(\underline{x}, \underline{\lambda}, \underline{\nu}).$$

Lemma 5.4.8. *The Lagrangian dual problem*

$$\begin{aligned} \text{maximize} \quad & g(\underline{\lambda}, \underline{\nu}) \\ \text{subject to} \quad & \underline{\lambda} \geq 0 \end{aligned}$$

has a unique maximum value $d^* \geq p^*$. This property is known as ***weak duality****.*

Proof. Since the Lagrangian dual function is the pointwise infimum of affine functions, it is always concave. Therefore, the function has a unique maximum value d^*. The upper bound on d^* follows from

$$\begin{aligned} g(\underline{\lambda}, \underline{\nu}) &= \inf_{\underline{x} \in \mathcal{D}} L(\underline{x}, \underline{\lambda}, \underline{\nu}) \\ &\leq \inf_{\underline{x} \in \mathcal{F}} L(\underline{x}, \underline{\lambda}, \underline{\nu}) \\ &= p^* + \sum_{i=1}^{m} \lambda_i f_i(\underline{x}) \\ &\leq p^*. \end{aligned}$$

□

The Lagrangian dual function can be $-\infty$ for a wide range of $(\underline{\lambda}, \underline{\nu})$. In this case, it makes sense to define the set of **implicit constraints** given by

$$\mathcal{C} = \left\{ (\underline{\lambda}, \underline{\nu}) | g(\underline{\lambda}, \underline{\nu}) > -\infty \right\}.$$

The points $(\underline{\lambda}, \underline{\nu}) \in \mathcal{C}$ are called **dual feasible**.

Definition 5.4.9. *If $d^* = p^*$, then one says that* ***strong duality*** *holds for the problem.*

Example 5.4.10. *The Lagrangian of a linear program (LP) in standard form is*

$$L(\underline{x}, \underline{\lambda}, \underline{\nu}) = \underline{c}^T \underline{x} + \underline{\nu}^T (A\underline{x} - \underline{b}) - \underline{\lambda}^T \underline{x}$$

and the Lagrangian dual function is given by

$$g(\underline{\lambda}, \underline{\nu}) = \inf_{\underline{x} \in \mathcal{D}} L(\underline{x}, \underline{\lambda}, \underline{\nu}) = \begin{cases} -\underline{b}^T \underline{\nu} & \text{if } A^T \underline{\nu} - \underline{\lambda} + \underline{c} = \underline{0} \\ -\infty & \text{otherwise}. \end{cases}$$

Solving the implicit constraint for $\underline{\lambda}$ and using the fact that $\underline{\lambda} \succeq \underline{0}$ gives the standard form of the LP dual problem

$$\begin{aligned} \text{maximize} \quad & -\underline{b}^T \underline{\nu} \\ \text{subject to} \quad & A^T \underline{\nu} + c \succeq \underline{0}. \end{aligned}$$

Strong duality for linear programs says that, if the original LP has an optimal solution (i.e., it is neither unbounded nor infeasible), then the dual LP has an optimal solution of the same value.

5.4.3 Convex Optimization

Definition 5.4.11. *An optimization problem in standard form is called* ***convex*** *if the function f_i is convex for $i = 0, 1, \ldots, m$, the function h_j is affine (i.e., $h_j(\underline{x}) = \underline{a}_j^T \underline{x} - b_j$) for $j = 1, 2, \ldots, p$, and $\mathcal{D} = \mathbb{R}^n$.*

Applying Theorem 5.3.4 to this setup shows that a standard convex optimization problem has a unique minimum value. Also, if the function f_0 is strictly convex, then the minimum value achieved uniquely. There are a number of stronger conditions that also imply strong duality for convex optimization problems. **Slater's condition** is stated below as a theorem and its proof can be found in [BV04, Sec. 5.3.2].

Theorem 5.4.12 (Slater's Condition)**.** *If a convex optimization problem has a point $\underline{x}_0$ where $f_i(\underline{x}_0) < 0$ for $i = 1, \ldots, m$ and $h_j(\underline{x}_0) = 0$ for $j = 1, \ldots, p$, then strong duality holds for the problem.*

Example 5.4.13. *For a channel with colored noise, the input distribution that maximizes the achievable information rate can be found by solving the convex optimiza-*

tion problem, known as water-filling, given by

$$\begin{aligned}\text{minimize} \quad & -\sum_{i=1}^{n} \log(x_i + \alpha_i) \\ \text{subject to} \quad & \sum_{i=1}^{n} x_i = P \\ & \underline{x} \succeq 0.\end{aligned}$$

Choosing $x_i = \frac{P}{n}$ *for* $i = 1, \ldots, n$ *gives a point that satisfies Slater's condition, so strong duality holds for this problem.*

Example 5.4.14. *For the water-filling problem, the Lagrangian can be written as*

$$L(\underline{x}, \underline{\lambda}, \nu) = -\sum_{i=1}^{n} \log(x_i + \alpha_i) - \sum_{i=1}^{m} \lambda_i x_i + \nu \left(-P + \sum_{i=1}^{n} x_i \right)$$

and the Lagrangian dual is given by $g(\underline{\lambda}, \nu) = \inf_{\underline{x} \in \mathbb{R}^n} L(\underline{x}, \underline{\lambda}, \nu)$.

If $\lambda_i < 0$, *then the Lagrangian tends to* $-\infty$ *as* $x_i \to -\infty$. *Thus, the system is implicitly constrained to have* $\lambda_i \geq 0$. *The first-order optimality conditions, for* $i = 1, 2, \ldots, n$, *are given by*

$$-\frac{1}{x_i + \alpha_i} - \lambda_i + \nu = 0.$$

Solving this for x_i *shows that* x_i *is increasing in* λ_i *(for* $\lambda_i \geq 0$*) and this implies that* $g(\underline{\lambda}, \nu)$ *is decreasing in* λ_i *(for* $\lambda_i \geq 0$ *and* $x_i \geq 0$*).*

Thus, the expression $\max_{\underline{\lambda} \geq 0} g(\underline{\lambda}, \nu)$ *is given by choosing the smallest non-negative* λ_i*'s for which* $x_i \geq 0$. *This implies that*

$$(x_i, \lambda_i) = \begin{cases} \left(\frac{1}{\nu} - \alpha_i, 0\right) & \text{if } \nu < \frac{1}{\alpha_i} \\ \left(0, \nu - \frac{1}{\alpha_i}\right) & \text{if } \nu \geq \frac{1}{\alpha_i}. \end{cases}$$

From this, the value of ν *can be determined by solving*

$$\sum_{i=1}^{n} x_i = \sum_{i=1}^{n} \max\left\{0, \frac{1}{\nu} - \alpha_i\right\} = P.$$

By strong duality, the optimal value of the dual problem equals the optimal value of the original problem. Finally, the problem can be easily solved for a range of P *values by sweeping through a range of* ν *values and computing* P *in terms of* ν.