EL3370 Mathematical Methods in Signals, Systems and Control

Topic 9: Differentiability and Optimization of Functionals

Cristian R. Rojas

Division of Decision and Control Systems
KTH Royal Institute of Technology

Outline

Outline

Differentiability

Inverse/Implicit Function Theorems

Calculus of Variations

Game Theory and the Minimax Theorem

Lagrangian Duality

Bonus Slides

Differentiability

Goal: *Generalize the notion of derivative to functionals on normed spaces.*

Definition. Let X, Y be normed spaces and $T: D \subseteq X \to Y$ (a possibly nonlinear transformation).
If, for $x \in D$, there exists a bounded linear operator $h \in X \mapsto \delta T(x;h) \in Y$ s.t.

$$\lim_{\|h\| \to 0} \frac{\|T(x+h) - T(x) - \delta T(x;h)\|}{\|h\|} = 0,$$

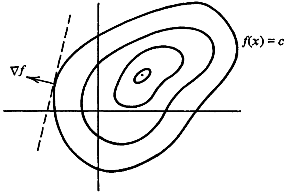

then T is *Fréchet differentiable* at x, and $\delta T(x;h)$ is the *Fréchet differential of T at x with increment h.*

If f is a functional on X, then $\delta f(x;h) = \left. \dfrac{d}{d\alpha} f(x + \alpha h) \right|_{\alpha=0}$.

Differentiability (cont.)

Examples

1. If $X = \mathbb{R}^n$ and $f(x) = f(x_1, \dots, x_n)$ is a functional having continuous partial derivatives with respect to each variable x_k, then

$$\delta f(x;h) = \sum_{k=1}^{n} \frac{\partial f}{\partial x_k} h_k.$$

2. Let $X = C[0,1]$ and $f(x) = \displaystyle\int_0^1 g(x(t),t)dt$ where g_x exists and is continuous with respect to x. Then $\delta f(x;h) = \displaystyle\frac{d}{d\alpha} \int_0^1 g(x(t)+\alpha h(t),t)dt\bigg|_{\alpha=0} = \int_0^1 g_x(x(t),t)h(t)dt.$

Properties

1. *If T has a Fréchet differential, it is unique.*
   Proof. If $\delta T(x;h)$, $\delta' T(x;h)$ are Fréchet differentials of T, and $\varepsilon > 0$, $\|\delta T(x;h) - \delta' T(x;h)\| \leqslant \|T(x+h) - T(x) - \delta T(x;h)\| + \|T(x+h) - T(x) - \delta' T(x;h)\| \leqslant \varepsilon \|h\|$ for h small. Thus, $\delta T(x;h) - \delta' T(x;h)$ is a bounded operator with norm 0, *i.e.*, $\delta T(x;h) = \delta' T(x;h)$. □

2. *If T is Fréchet differentiable at $x \in D$, where D is open, then T is continuous at x.*
   Proof. Given $\varepsilon > 0$, there is a $\delta > 0$ s.t. $\|T(x+h) - T(x) - \delta T(x;h)\| \leqslant \varepsilon \|h\|$ whenever $\|h\| < \delta$, *i.e.*, $\|T(x+h) - T(x)\| < \varepsilon \|h\| + \|\delta T(x;h)\| \leqslant (\varepsilon + M)\|h\|$, where $M = \|\delta T(x;\cdot)\|$, so T is continuous at x. □

If $T: D \subseteq X \to Y$ is Fréchet differentiable throughout D, then the Fréchet differential is of the form $\delta T(x;h) = T'(x)h$, where $T'(x) \in \mathscr{L}(X,Y)$ is the *Fréchet derivative of T at x*.

Also, if $x \mapsto T'(x)$ is continuous in some open $S \subseteq D$, then T is *continuously Fréchet differentiable in S*.

If f is a functional in D, so that $\delta f(x;h) = f'(x)h$, $f'(x) \in X^*$ is the *gradient of f at x*.

Much of the theory for ordinary derivatives extends to Fréchet derivatives:

Properties

1. *(Chain rule). Let $S : D \subseteq X \to E \subseteq Y$ and $P : E \to Z$ be Fréchet differentiable at $x \in D$ and $y = S(x) \in E$, respectively, where X, Y, Z are normed spaces and D, E are open sets. Then $T = P \circ S$ is Fréchet differentiable at x, and $T'(x) = P'(y)S'(x)$.*

   Proof. If $x, x+h \in D$, then $T(x+h) - T(x) = P[S(x+h)] - P[S(x)] = P(y+g) - P(y)$, where $g = S(x+h) - S(x)$. Now, $\|P(y+g) - P(y) - P'(y)g\| = o(\|g\|)$, $\|g - S'(x)h\| = o(\|h\|)$ and $\|g\| = O(\|h\|)$, so $\|T(x+h) - T(x) - P'(y)S'(x)h\| = o(\|h\|) + o(\|g\|) = o(\|h\|)$. Thus, $T'(x)h = P'(y)S'(x)h$. $\square$

Properties (cont.)

2. *(Mean value theorem). Let T be Fréchet differentiable on an open domain D, and $x \in D$. Suppose that $x + th \in D$ for all $t \in [0,1]$. Then $\|T(x+h) - T(x)\| \leqslant \|h\| \sup_{0<t<1} \|T'(x+th)\|$.*

   Fix $y^* \in D^*$, $\|y^*\| = 1$, and let $\phi(t) := \langle T(x+th), y^* \rangle$ ($t \in [0,1]$), which is differentiable, with $\phi'(t) = \langle T'(x+th)h, y^* \rangle$. Let $\gamma(t) = \phi(t) - (1-t)\phi(0) - t\phi(1)$, so $\gamma(0) = \gamma(1) = 0$ and $\gamma'(t) = \phi'(t) + \phi(0) - \phi(1)$. If $\gamma = 0$, then $\gamma' = 0$; if not, there is a $\tau \in (0,1)$ s.t., *e.g.*, $\gamma(\tau) > 0$, so there is an $s \in (0,1)$ s.t. $\gamma(s) = \max_{t \in [0,1]} \gamma(t)$. Now, $\gamma(s+h) - \gamma(s) \leqslant 0$ whenever $0 \leqslant s + h \leqslant 1$, so $\gamma'(s) = 0$, and $|\phi(1) - \phi(0)| = |\phi'(s)| \leqslant \sup_{0<t<1} |\phi'(t)| \leqslant \|h\| \sup_{0<t<1} \|T'(x+th)\|$. Also, $|\phi(1) - \phi(0)| = |\langle T(x+h) - T(x), y^* \rangle|$, so taking the sup over $\|y^*\| = 1$ yields the result. $\square$

Extrema

The minima/maxima of a functional can be found by setting its Fréchet derivative to zero!

Definition. $x_0 \in \Omega$ is a *local minimum* of $f : \Omega \subseteq X \to \mathbb{R}$ if there is a nbd B of x where $f(x_0) \leqslant f(x)$ on $\Omega \cap B$, and a *strict local minimum* if $f(x_0) < f(x)$ for all $x \in \Omega \cap B \setminus \{x_0\}$.

Theorem. If $f : X \to \mathbb{R}$ is Fréchet differentiable, then a necessary condition for f to have a local minimum/maximum at $x_0 \in X$ is that $\delta f(x_0; h) = 0$ for all $h \in X$.

Proof. If $\delta f(x_0; h) \neq 0$, pick h_0 s.t. $\|h_0\| = 1$ and $\delta f(x_0; h_0) > 0$. As $h \to 0$, $|f(x_0 + h) - f(x_0) - \delta f(x_0; h)|/\|h\| \to 0$, so given $\varepsilon \in (0, \delta f(x_0; h_0))$ there is a $\gamma > 0$ s.t. $f(x_0 + \gamma h_0) > f(x_0) + \delta f(x_0; \gamma h_0) - \varepsilon\gamma > f(x_0)$, while $f(x_0 - \gamma h_0) < f(x_0) - \delta f(x_0; \gamma h_0) + \varepsilon\gamma < f(x_0)$, so x_0 is not a local minimum/maximum. $\square$

A generalization of this result to constrained optimization is:

Theorem. If x_0 minimizes f on the convex set $\Omega \subseteq X$, and f is Fréchet differentiable at x_0, then $\delta f(x_0; x - x_0) \geqslant 0$ for all $x \in \Omega$.

Proof. For $x \in \Omega$, let $h = x - x_0$ and note that $x_0 + \alpha h \in \Omega$ $(0 \leqslant \alpha \leqslant 1)$ since Ω is convex. The rest of the proof is similar to the previous one. $\square$

Outline

Differentiability

Inverse/Implicit Function Theorems

Calculus of Variations

Game Theory and the Minimax Theorem

Lagrangian Duality

Bonus Slides

Inverse/Implicit Function Theorems

The inverse and implicit function theorems are fundamental to many fields, and constitute the analytical backbone of differential geometry, essential to nonlinear system theory.

Theorem (Inverse Function Theorem)
Let X, Y be Banach spaces, and $x_0 \in X$. Assume that $T: X \to Y$ is continuously Fréchet differentiable in a nbd of x_0, and that $T'(x_0)$ is invertible. Then, there is a nbd U of x_0 s.t. T is invertible in U, and both T and T^{-1} are continuous. Furthermore, T^{-1} is continuously Fréchet differentiable in $T(U)$, with derivative $[T'(T^{-1}(y))]^{-1}$ ($y \in T(U)$).

Proof.

(1) *Invertibility*: Since $T'(x_0)$ is invertible, by translation and multiplying by a linear map, assume w.l.o.g. that $x_0 = 0$, $T(x_0) = 0$ and $T'(x_0) = I$. Consider $y \mapsto T_y(x) = x - T(x) + y$ for $y \in X$; note that a fixed point of T_y is precisely an x s.t. $T(x) = y$. Define the ball $\overline{B_R} := \{x \in X : \|x\| \leqslant R\}$, which is complete. Let $F(x) = T(x) - x$. By the mean value theorem, $\|F(x) - F(x')\| \leqslant \sup_{z \in \overline{B_R}} \|F'(z)\| \cdot \|x - x'\|$ for all $x, x' \in \overline{B_R}$, and since $F'(0) = 0$, given a fixed $\varepsilon \in (0,1)$, if $R > 0$ is small enough, $\|F(x) - F(x')\| \leqslant \varepsilon \|x - x'\|$.

Proof (cont.)

Suppose $\|y\| \leqslant R(1-\varepsilon)$. Note that, if $x \in \overline{B_R}$, $\|T_y(x)\| \leqslant \|F(x)\| + \|y\| \leqslant \varepsilon\|x\| + R(1-\varepsilon) \leqslant R$, so $T_y(\overline{B_R}) \subseteq \overline{B_R}$, and for $x, x' \in \overline{B_R}$, $\|T_y(x) - T_y(x')\| \leqslant \|F(x) - F(x')\| \leqslant \varepsilon\|x - x'\|$, so T_y is a contraction. By the Banach fixed point theorem (Topic 4), T_y has a unique fixed point, *i.e.*, if $\|y\|$ is small enough, there is a unique $x \in \overline{B_R}$ s.t. $T(x) = y$, so $T^{-1}: \overline{B_{R(1-\varepsilon)}} \to \overline{B_R}$ exists.

(2) *Continuity*: Since T is Fréchet differentiable in $\overline{B_R}$, it is continuous there. For $y, y_0 \in \overline{B_{R(1-\varepsilon)}}$, $\|T_y(x) - T_{y_0}(x)\| = \|y - y_0\| \to 0$ as $y \to y_0$, so by the last part of the Banach fixed point theorem, T^{-1} is continuous.

(3) *Continuous differentiability*: Consider a nbd $V \subseteq \overline{B_R}$ of 0 where T' is invertible. Let $W = T(V)$, $y_0, y \in W$ and $x_0 = T^{-1}(y_0)$, $x = T^{-1}(y)$. Then,

$$\begin{aligned}\frac{\|T^{-1}(y) - T^{-1}(y_0) - [T'(x_0)]^{-1}(y - y_0)\|}{\|y - y_0\|} &= \frac{\|x - x_0 - [T'(x_0)]^{-1}(T(x) - T(x_0))\|}{\|T(x) - T(x_0)\|} \\ &\leqslant \|[T'(x_0)]^{-1}\| \left(\frac{\|T(x) - T(x_0) - [T'(x_0)](x - x_0)\|}{\|x - x_0\|}\right)\left(\frac{\|x - x_0\|}{\|T(x) - T(x_0)\|}\right). \quad (*)\end{aligned}$$

Proof (cont.)

The 2nd factor tends to 0 as $x \to x_0$, while for the 3rd factor:

$$\begin{aligned}\liminf_{x\to x_0} \frac{\|\, T(x) - T(x_0)\|}{\|x - x_0\|} &\geqslant \liminf_{x\to x_0} \left| \frac{\|T'(x_0)[x - x_0]\|}{\|x - x_0\|} - \frac{\|T(x) - T(x_0) - T'(x_0)[x - x_0]\|}{\|x - x_0\|} \right| \\ &= \liminf_{x\to x_0} \frac{\|T'(x_0)[x - x_0]\|}{\|x - x_0\|} \geqslant \frac{1}{\|[T'(x_0)]^{-1}\|} > 0.\end{aligned}$$

Hence, the left hand side of $(*)$ tends to 0, and $T^{-1}(y_0)$ has Fréchet derivative $[T'(x_0)]^{-1}$. □

Theorem (Implicit Function Theorem)
Let X, Y, Z be Banach spaces, $A \subseteq X \times Y$ open, and $f: A \to Z$ continuously Fréchet differentiable, with derivative $[f_x \ \ f_y]$. Let $(x_0, y_0) \in A$ be s.t. $f(x_0, y_0) = 0$, and assume that $f_y(x_0, y_0)$ is invertible. Then, there are open sets $W \subseteq X$ and $V \subseteq A$ s.t $x_0 \in W$, $(x_0, y_0) \in V$, and a $g: W \to Y$ Fréchet differentiable at x_0 s.t. $(x, g(x)) \in V$ and $f(x, g(x)) = 0$ for all $x \in W$. Moreover, $g'(x_0) = -[f_y(x_0, y_0)]^{-1} f_x(x_0, y_0)$.

Proof. Define the continuously differentiable function $F: A \to X \times Z$ by $F(x, y) = (x, f(x, y))$. Note that $F(x_0, y_0) = (x_0, 0)$ and

$$F'(x_0, y_0) = \begin{bmatrix} I & 0 \\ f_x(x_0, y_0) & f_y(x_0, y_0) \end{bmatrix}, \qquad [F'(x_0, y_0)]^{-1} = \begin{bmatrix} I & 0 \\ -[f_y(x_0, y_0)]^{-1} f_x(x_0, y_0) & [f_y(x_0, y_0)]^{-1} \end{bmatrix},$$

i.e., $F'(x_0, y_0)$ is invertible. By the inverse function theorem, there is an open $V \subseteq A$ where F is invertible and F^{-1} is continuously differentiable. Let $\pi_Y : X \times Y \to Y$ be the projection of $X \times Y$ onto Y, *i.e.*, $\pi_Y(x, y) = y$ for all $(x, y) \in X \times Y$. The function $g: W \to Y$ given by $g(x) = \pi_Y(F^{-1}(x, 0))$, where $W = \{x \in X : (x, 0) \in F(V)\}$, satisfies the conditions of the theorem. □

Application to initial-value problems

Consider the initial-value problem

$$\frac{dx(t)}{dt} = f(x,t), \quad t \in [a,b]$$
$$x(a) = \xi \in \mathbb{R}^n,$$

where f is continuously differentiable, and $x \in C([a,b],\mathbb{R}^n)$.

We want to study the dependence of x on ξ. To this end, define the function $\Phi\colon C([a,b],\mathbb{R}^n) \times \mathbb{R}^n \to C([a,b],\mathbb{R}^n)$ as

$$\Phi(x,\xi)(t) = x(t) - \xi - \int_a^t f(x(s),s)ds, \quad t \in [a,b].$$

Notice that x solves the initial-value problem iff $\Phi(x,\xi) = 0$. Now, Φ is continuously differentiable, and it satisfies the conditions of the implicit function theorem (*check this!*), which implies that x depends on ξ in a differentiable manner!

Outline

Classical problem: find a function x on $[t_1, t_2]$ that minimizes $J = \int_{t_1}^{t_2} f[x(t), \dot{x}(t), t] dt$.

Assume that x belongs to the space $D[t_1, t_2]$ of real-valued continuously differentiable functions on $[t_1, t_2]$, with norm $\|x\| = \max_{t_1 \leqslant t \leqslant t_2} |x(t)| + \max_{t_1 \leqslant t \leqslant t_2} |\dot{x}(t)|$.
Also, the end points $x(t_1)$ and $x(t_2)$ are assumed fixed.

If $D_h[t_1, t_2]$ is the subspace consisting of those $x \in D[t_1, t_2]$ s.t. $x(t_1) = x(t_2) = 0$, then the necessary condition for the minimization of J is

$$\delta J(x; h) = 0, \quad \text{for all } h \in D_h[t_1, t_2].$$

Calculus of Variations (cont.)

We have

$$\begin{aligned}\delta J(x;h) &= \frac{d}{d\alpha}\int_{t_1}^{t_2} f(x+\alpha h, \dot{x}+\alpha\dot{h}, t)dt\bigg|_{\alpha=0} \\ &= \int_{t_1}^{t_2} f_x(x,\dot{x},t)h(t)dt + \int_{t_1}^{t_2} f_{\dot{x}}(x,\dot{x},t)\dot{h}(t)dt \\ &= \int_{t_1}^{t_2}\left[f_x(x,\dot{x},t) - \frac{d}{dt}f_{\dot{x}}(x,\dot{x},t)\right]h(t)dt.\end{aligned}$$

(integration by parts, assuming $\frac{d}{dt}f_{\dot{x}}(x,\dot{x},t)$ is continuous in t)

Lemma (Fundamental lemma of calculus of variations)
If $\alpha \in C[t_1,t_2]$, and $\int_{t_1}^{t_2}\alpha(t)h(t)dt = 0$ for every $h \in D_h[t_1,t_2]$, then $\alpha = 0$.

Proof. If, say, $\alpha(t) > 0$ for some $t \in (t_1,t_2)$, there is an interval (τ_1,τ_2) where α is strictly positive. Pick $h(t) = (t-\tau_1)^2(t-\tau_2)^2$ for $t \in (\tau_1,\tau_2)$ and $h(t) = 0$ otherwise. This gives $\int_{t_1}^{t_2}\alpha(t)h(t)dt > 0$. $\square$

Using this result we obtain

$$\delta J(x;h) = 0 \text{ for all } h \in D_h[t_1,t_2] \quad \Leftrightarrow \quad f_x(x,\dot{x},t) - \frac{d}{dt}f_{\dot{x}}(x,\dot{x},t) = 0. \quad \textit{(Euler-Lagrange equation)}$$

Example (minimum arc length)
Problem: Given $(t_1, x(t_1)), (t_2, x(t_2))$, determine curve of minimum length connecting them.

Notice that the distance between points $(t, x(t))$ and $(t + \Delta t, x(t + \Delta t))$ is

$$\sqrt{(x(t + \Delta t) - x(t))^2 + \Delta t^2} = \sqrt{(\dot{x}(t)\Delta t + o(\Delta t))^2 + \Delta t^2} = \sqrt{1 + \dot{x}^2(t)}\,\Delta t + o(\Delta t),$$

hence the total arc length, by integration, is: $J = \int_{t_1}^{t_2} \sqrt{1 + \dot{x}^2(t)}dt.$

Using the Euler-Lagrange equation, we obtain

$$\frac{d}{dt}\frac{\partial}{\partial \dot{x}}\sqrt{1 + \dot{x}^2} = 0$$

or $\dot{x}(t) = \text{constant}$. Thus, the extremizing arc is the straight line connecting these points.

Outline

Game Theory and the Minimax Theorem

Two-Person Zero-Sum Games

Consider a problem with two players: I and II. If player I chooses a *strategy* $x \in X$, and player II chooses a *strategy* $y \in Y$, then I *gains*, and II *loses*, an amount (*payoff*) $J(x,y)$. Each player wants to maximize its payoff.

Example: Matching pennies

Player II

Player I		y_1	y_2
	x_1	-1	1
	x_2	1	-1

Player I wants to maximize $\min_{y \in Y} J(x,y)$ wrt x.
Player II wants to minimize $\max_{x \in X} J(x,y)$ wrt y.

If $V_* = \max_{x \in X} \min_{y \in Y} J(x,y)$ and $V^* = \min_{y \in Y} \max_{x \in X} J(x,y)$, and $V^* = V_*$, $V = V^* = V_*$ is the *value* of the game.

Not every game has a value!

Mixed Strategies

Instead of choosing a particular strategy, each player can choose a *mixed/randomized strategy*, *i.e.*, a probability distribution over its strategy space X or Y: $p_x(x)$, $p_y(y)$ (assuming that X and Y are finite).

The values of the game are

$$V_* = \max_{p_x} \min_{p_y} \sum_{x \in X} \sum_{y \in Y} J(x,y) p_x(x) p_y(y),$$

$$V^* = \min_{p_y} \max_{p_x} \sum_{x \in X} \sum_{y \in Y} J(x,y) p_x(x) p_y(y).$$

The fundamental (minimax) theorem of game theory states that $V_* = V^*$.

Proof of Minimax Theorem
We need to establish, equivalently, that for any matrix $A \in \mathbb{R}^{m \times n}$

$$V_* := \max_{\substack{x \in (\mathbb{R}_0^+)^n \\ x_1+\cdots+x_n=1}} \min_{\substack{y \in (\mathbb{R}_0^+)^m \\ y_1+\cdots+y_m=1}} x^T A y = \min_{\substack{y \in (\mathbb{R}_0^+)^m \\ y_1+\cdots+y_m=1}} \max_{\substack{x \in (\mathbb{R}_0^+)^n \\ x_1+\cdots+x_n=1}} x^T A y =: V^*.$$

First notice that, for every x, y:

$$\min_{\substack{y' \in (\mathbb{R}_0^+)^m \\ y'_1+\cdots+y'_m=1}} x^T A y' \leqslant x^T A y \leqslant \max_{\substack{x' \in (\mathbb{R}_0^+)^n \\ x'_1+\cdots+x'_n=1}} x'^T A y.$$

so taking max wrt x and min wrt y gives $V_* \leqslant V^*$.

We need to show that $V_* \geqslant V^*$, by showing that there is an x_0 s.t. $\min_{\substack{y \in (\mathbb{R}_0^+)^m \\ y_1+\cdots+y_m=1}} x_0^T A y = V^*$.

Game Theory and the Minimax Theorem (cont.)

Reformulation as an S-game

To gain geometric insight, we can simplify the problem by defining the *risk set*

$$S := \{Ay \in \mathbb{R}^n : y \in (\mathbb{R}_0^+)^m,\ y_1 + \cdots + y_m = 1\}$$

so $\displaystyle \min_{\substack{y \in (\mathbb{R}_0^+)^m \\ y_1 + \cdots + y_m = 1}} x^T A y = \min_{s \in S} x^T s.$

Example

	y_1	y_2	y_3	y_4
x_1	2	2	0	0.5
x_2	0	1	2	1

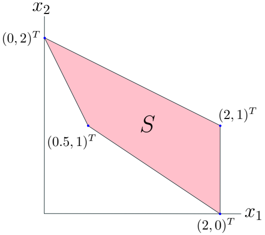

S is the *convex hull* of the columns of A.

Game Theory and the Minimax Theorem (cont.)

Back to the proof...

A minimax strategy for Player II, *i.e.*, an y_0 s.t. $\max_{\substack{x \in (\mathbb{R}_0^+)^n \\ x_1+\cdots+x_n=1}} x^T A y_0 = \min_{\substack{y \in (\mathbb{R}_0^+)^m \\ y_1+\cdots+y_m=1}} \max_{\substack{x \in (\mathbb{R}_0^+)^n \\ x_1+\cdots+x_n=1}} x^T A y,$

corresponds to an $s_0 = A y_0 \in S$ of minimum $s_{\max} := \max\{s_1, \ldots, s_n\}$.

Let $Q_\alpha := \{s \in \mathbb{R}^n : s_{\max} \leqslant \alpha\}$. Then

$$V^* = \inf\{\alpha \in \mathbb{R} : Q_\alpha \cap S \neq \emptyset\}.$$

To find an x_0 s.t. $\min_{s \in S} x_0^T s = V^*$, we can use the separating hyperplane theorem to determine a hyperplane (given by $\bar{x}$) separating Q_{V^*} and S: $H = \{s \in S : \bar{x}^T s = V^*\}$.

($\bar{x}$ has been scaled so that $\sum_j \bar{x}_j = 1$, since H contains the vertex $s^* = (V^*, \ldots, V^*)$ of Q_{V^*}, so $\bar{x}^T s_0 = \bar{x}^T s^* = V^* \sum_j \bar{x}_j = V^*$, and $\bar{x}^T s \leqslant V^*$ for all $s \in Q_\alpha$ implies, by letting $s_j \to -\infty$, that $\bar{x}_j \geqslant 0$ for all j).

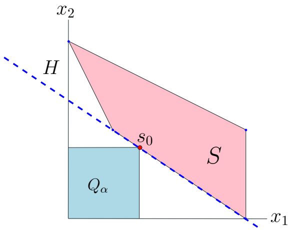

Then we can choose $x_0 = \bar{x}$! This proves the minimax theorem. $\square$

Outline

Lagrangian Duality

Given a convex optimization problem in a normed space, our goal is to derive its (Lagrangian) dual. To formulate such a problem, we need to define an order relation:

Definitions

- A set C in a real vector space V is a *cone* if $x \in C$ implies that $\alpha x \in C$ for every $\alpha \geqslant 0$.
- Given a convex cone P in V (*positive cone*), we say that $x \geqslant y$ ($x, y \in V$) when $x - y \in P$.
- If V is a normed space with closed positive cone P, $x > 0$ means that $x \in \operatorname{int} P$.
- Given the positive cone $P \subseteq V$, $P^{\oplus} := \{x^* \in V^* : x^*(x) \geqslant 0 \text{ for all } x \in P\}$ is the *positive cone in* V^*. By Hahn-Banach, if P is closed and $x \in V$, then $x^*(x) \geqslant 0$ for all $x^* \geqslant 0$ implies that $x \geqslant 0$.
- If X, Y are real vector spaces, $C \subseteq X$ is convex, and P is the positive cone of Y, a function $f : C \to Y$ is *convex* if $f(\alpha x + (1-\alpha)y) \leqslant \alpha f(x) + (1-\alpha) f(y)$ for all $x, y \in X$, $\alpha \in [0, 1]$.

Given a vector space X and a normed space Y, let Ω be a convex subset of X, and P be the (closed) positive cone of Y. Also, let $f : \Omega \to \mathbb{R}$ and $G : \Omega \to Y$ be convex functions.

Consider the convex optimization problem

$$\begin{aligned} \min_{x \in X} \quad & f(x) \\ \text{s.t.} \quad & x \in \Omega, \; G(x) \leqslant 0. \end{aligned}$$

To analyze this convex optimization problem, we need to introduce a special function:

Definition

Let $\Gamma = \{y \in Y : \text{there exists an } x \in \Omega \text{ s.t. } G(x) \leqslant y\}$; this set is convex (*why?*).
On Γ, the *primal function* $\omega \colon \Omega \to \mathbb{R}$ is given by $\omega(y) := \inf\{f(x) \colon x \in \Omega,\, G(x) \leqslant y\}$.
Notice that the original optimization problem corresponds to finding $\omega(0)$.

Properties

(1) ω is convex.

Proof

$$\begin{aligned}
&\omega(\alpha y_1 + (1-\alpha)y_2) \\
&= \inf\{f(x) \colon x \in \Omega,\, G(x) \leqslant \alpha y_1 + (1-\alpha)y_2\} \\
&\leqslant \inf\{f(\alpha x_1 + (1-\alpha)x_2) \colon x_1, x_2 \in \Omega,\, G(x_1) \leqslant y_1,\, G(x_2) \leqslant y_2\} \\
&\leqslant \alpha \inf\{f(x) \colon x \in \Omega,\, G(x) \leqslant y_1\} \\
&\quad + (1-\alpha) \inf\{f(x) \colon x \in \Omega,\, G(x) \leqslant y_2\} \\
&= \alpha\omega(y_1) + (1-\alpha)\omega(y_2).
\end{aligned}$$

□

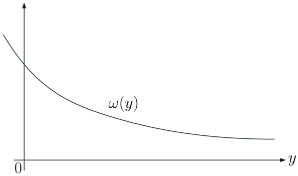

(2) ω is *non-increasing*: if $y_1 \leqslant y_2$ then $\omega(y_1) \geqslant \omega(y_2)$.

Proof. Direct. □

Duality theory of convex programming is based on the observation that, since ω is convex, its *epigraph* (*i.e.*, the area above the curve of ω in $\Gamma \times \mathbb{R}$) is convex, so it has a supporting hyperplane passing through the point $(0, \omega(0))$. To develop this idea, consider the normed space $Y \times \mathbb{R}$ with the norm $\|(y,r)\| = \|y\| + |r|$ for $y \in Y$ and $r \in \mathbb{R}$.

Theorem

Assume that P has non-empty interior, and that there exists an $x_1 \in \Omega$ s.t. $G(x_1) < 0$ (*i.e.*, $-G(x_1)$ is an interior point of P). Let

$$\mu_0 = \inf\{f(x) \colon x \in \Omega, G(x) \leqslant 0\}, \tag{$*$}$$

and assume μ_0 is finite. Then, there exists a $y_0^* \in P^\oplus$ s.t.

$$\mu_0 = \inf\{f(x) + \langle G(x), y_0^* \rangle \colon x \in \Omega\}. \tag{$**$}$$

Furthermore, if the infimum in $(*)$ is achieved by some $x_0 \in \Omega$, $G(x_0) \leqslant 0$, then the infimum in $(**)$ is also achieved by x_0, and $\langle G(x_0), y_0^* \rangle = 0$.

Lagrangian Duality (cont.)

Proof. On $Y \times \mathbb{R}$, define the sets

$$A := \{(y,r)\colon\ y \geqslant G(x),\ r \geqslant f(x),\ \text{for some } x \in \Omega\}, \qquad \text{(epigraph of } f\text{)}$$
$$B := \{(y,r)\colon\ y \leqslant 0,\ r \leqslant \mu_0\}.$$

Since f, G are convex, so are the sets A, B. By the definition of μ_0, $A \cap \operatorname{int} B = \emptyset$. Also, since P has an interior point, B has a non-empty interior (*why?*). Then, by the separating hyperplane theorem, there is a non-zero $w_0^* = (y_0^*, r_0) \in (Y \times \mathbb{R})^*$ s.t.

$$\langle y_1, y_0^* \rangle + r_0 r_1 \geqslant \langle y_2, y_0^* \rangle + r_0 r_2, \qquad \text{for all } (y_1, r_1) \in A,\ (y_2, r_2) \in B.$$

From the nature of B, it follows that $y_0^* \geqslant 0$ and $r_0 \geqslant 0$. Since $(0, \mu_0) \in B$, we have that $\langle y, y_0^* \rangle + r_0 r \geqslant r_0 \mu_0$ for all $(y,r) \in A$; if $r_0 = 0$, then in particular $y_0^* \neq 0$ and $\langle G(x_1), y_0^* \rangle \geqslant 0$, but since $-G(x_1) > 0$ and $y_0^* \geqslant 0$, we would have that $\langle G(x_1), y_0^* \rangle < 0$ (we know that $\langle G(x_1), y_0^* \rangle \leqslant 0$; now, there exists a $y \in Y$ s.t. $\langle y, y_0^* \rangle > 0$, so $G(x_1) + \varepsilon y < 0$ for some $\varepsilon > 0$, thus if $\langle G(x_1), y_0^* \rangle = 0$ we would have $\langle G(x_1) + \varepsilon y, y_0^* \rangle > 0$, a contradiction). Therefore, $r_0 > 0$, and we can assume w.l.o.g. that $r_0 = 1$.
Since $(0, \mu_0) \in A \cap B$, $\mu_0 = \inf\{\langle y, y_0^* \rangle + r \colon (y,r) \in A\} = \inf\{f(x) + \langle G(x), y_0^* \rangle \colon x \in \Omega\} \leqslant \inf\{f(x)\colon x \in \Omega, G(x) \leqslant 0\}$ $= \mu_0$, which establishes the first part of the theorem. Now, if there is an $x_0 \in \Omega$ s.t. $G(x_0) \leqslant 0$ and $f(x_0)$ $= \mu_0$, then $\mu_0 \leqslant f(x_0) + \langle G(x_0), y_0^* \rangle \leqslant f(x_0) = \mu_0$, so $\langle G(x_0), y_0^* \rangle = 0$. □

The expression $L(x, y^*) = f(x) + \langle G(x), y^* \rangle$, for $x \in \Omega$, $y^* \in P^\oplus$, is the *Lagrangian* of the optimization problem.

Corollary (Lagrangian Dual). Under the conditions of the theorem,

$$\sup_{y' \in P^\oplus} \mathscr{L}(y^*) := \inf\{f(x) + \langle G(x), y^* \rangle : x \in \Omega\} = \mu_0,$$

and the supremum is achieved by some $y_0^* \in P^\oplus$.

Proof. The theorem established the existence of a y_0^* s.t. $\mathscr{L}(y^*) = \mu_0$, while for all $y^* \in P^\oplus$, $\mathscr{L}(y^*) = \inf_{x \in \Omega}(f(x) + \langle G(x), y^* \rangle) \leqslant \inf_{x \in \Omega, G(x) \leqslant 0}(f(x) + \langle G(x), y^* \rangle) \leqslant \inf_{x \in \Omega, G(x) \leqslant 0} f(x) = \mu_0$. □

The dual problem can provide useful information about the primal (original) problem, since their solutions are linked via the complementarity condition $\langle G(x_0), z_0^* \rangle = 0$. Also, the dual problem always has a solution, so it may be easier to analyze than the primal.

Remark. If f is non-convex, ω may be non-convex, and the optimal cost of the dual problem provides only a *lower bound* on the optimal cost of the original problem.

Examples

(1) **Linear programming.** Let $A \in \mathbb{R}^{m\times n}$, $b \in \mathbb{R}^n$ and $c \in \mathbb{R}^m$, and consider the problem

$$\begin{aligned} \mu_0 = \min_{x\in\mathbb{R}^n} \quad & b^T x \\ \text{s.t.} \quad & Ax \geqslant c, \quad x \geqslant 0. \end{aligned}$$

Assume there is an $x > 0$ with $Ax > c$. Letting $f(x) = b^T x$, $G(x) = c - Ax$ and $\Omega = P = \{x \colon x_j \geqslant 0 \text{ for all } j\}$, the corollary yields, for $y \in P^\oplus = P$,

$$\mathscr{L}(y) = \inf\{b^T x + y^T(c - Ax) \colon x \geqslant 0\} = \inf\{(b - A^T y)^T x + y^T c \colon x \geqslant 0\} = \begin{cases} y^T c, & \text{if } b \geqslant A^T y \\ -\infty, & \text{otherwise,} \end{cases}$$

so the Lagrangian dual, corresponding to the standard dual linear program, is

$$\begin{aligned} \mu_0 = \max_{y\in\mathbb{R}^m} \quad & c^T y \\ \text{s.t.} \quad & A^T y \leqslant b, \quad y \geqslant 0. \end{aligned}$$

Lagrangian Duality (cont.)

Examples (cont.)

(2) **Optimal control.** Consider the system $\dot{x}(t) = Ax(t) + bu(t)$, where $x(t) \in \mathbb{R}^n$, $A \in \mathbb{R}^{n\times n}$, $b \in \mathbb{R}^n$ and $u(t) \in \mathbb{R}$. Given $x(t_0)$, the goal is to find an input u on $[t_0, t_1]$ which minimizes

$$J(u) = \int_{t_0}^{t_1} u^2(t)dt,$$

while satisfying $x(t_1) \geqslant c$, where $c \in \mathbb{R}^n$. The solution of the system is

$$x(t_1) = e^{A(t_1-t_0)}x(t_0) + Ku, \qquad Ku := \int_{t_0}^{t_1} e^{A(t_1-t)}bu(t)dt,$$

so problem corresponds to minimizing $J(u)$ subject to $Ku \geqslant c - e^{A(t_1-t_0)}x(t_0) =: d$.

Assuming that $u \in L_2[t_0, t_1]$, the corollary gives the dual problem

$$\begin{aligned}\max_{y\geqslant 0} \inf_{u\in L_2[t_0,t_1]} [J(u) + y^T(d - Ku)] &= \max_{y\geqslant 0} \inf_{u\in L_2[t_0,t_1]} \int_{t_0}^{t_1} [u^2(t) - y^T e^{A(t_1-t)}bu(t)]dt + y^T d \\ &= \max_{y\geqslant 0} y^T Q y + y^T d,\end{aligned}$$

where $Q := -(1/4)\int_{t_0}^{t_1} e^{A(t_1-t)}bb^T e^{A^T(t_1-t)}dt$. This is a simple finite-dimensional problem, and its solution, y_{opt}, yields $u_{\text{opt}}(t) = (1/2)y_{\text{opt}}^T e^{A(t_1-t)}b$.

Next Topic

Application to H_∞ Control Theory

Outline

Bonus: Equality Constrained Optimization

Problem

$$\begin{array}{ll} \min\limits_{x\in\Omega} & f(x) \\ \text{s.t.} & g_j(x) = 0, \quad j = 1,\dots,n, \end{array}$$

where $\Omega \subseteq X$ and $f, g_1, \dots, g_n$ are Fréchet differentiable on X.

Theorem 1. Let $x_0 \in \Omega$ be a local minimum of f on the set of all $x \in \Omega$ s.t. $g_j(x) = 0$, $j = 1,\dots,n$, and assume that the functionals $\delta g_1(x_0;\cdot),\dots,\delta g_n(x_0;\cdot)$ are l.i. (*i.e.*, x_0 is a *regular point*). Then,

$$\delta f(x_0;h) = 0 \text{ for all } h \text{ s.t. } \delta g_j(x_0;h) = 0 \text{ for all } j = 1,\dots,n.$$

Bonus: Equality Constrained Optimization (cont.)

Proof
First, notice that there exist vectors $y_1, \dots, y_n \in X$ s.t. the matrix $M \in \mathbb{R}^{n\times n}$, $M_{jk} = \delta g_j(x_0; y_k)$, is non-singular. To see this, consider the linear mapping $G\colon X \to \mathbb{R}^n$, $[G(y)]_j = \delta g_j(x_0; y)$. The range of G is a linear subspace of $\mathbb{R}^n$; if $\dim \mathscr{R}(G) < n$, there would exist a $\lambda \in \mathbb{R}^n \setminus \{0\}$ s.t. $\lambda^T G(y) = 0$ for all $y \in X$, *i.e.*, $\{\delta g_j(x_0; \cdot)\}$ would be l.i. Therefore, in particular there exist vectors $y_1, \dots, y_n \in X$ s.t. $G(y_j) = e_j$, so $M = I$, which is non-singular.

Fix $h \in X$ s.t. $\delta g_j(x_0; h) = 0$ for all $j = 1, \dots, n$, and consider the set of equations $g_j(x_0 + \alpha h + \sum_{k=1}^n \beta_k y_k) = 0$, $k = 1, \dots, n$, in $\alpha, \beta_1, \dots, \beta_n$. The Jacobian of this system, $[\partial g_j / \partial \beta_k]_{\alpha=\beta_k=0} = M$, is non-singular. Therefore, by the implicit function theorem (in $\mathbb{R}^n$), there exists a continuous function $\beta\colon U \subseteq \mathbb{R} \to \mathbb{R}^n$ in a nbd U of 0 s.t. $\beta(0) = 0$ and

$$\begin{aligned} 0 &= g_j\left(x_0 + \alpha h + \textstyle\sum_{k=1}^n \beta_k(\alpha) y_k\right) \\ &= \underbrace{g_j(x_0) + \alpha \delta g_j(x_0; h)}_{=0} + \underbrace{\delta g_j\left(x_0; \textstyle\sum_{k=1}^n \beta_k(\alpha) y_k\right)}_{=M\beta(\alpha)} + o(\alpha) + o\left(\left\|\textstyle\sum_{k=1}^n \beta_k(\alpha) y_k\right\|\right). \end{aligned}$$

Proof (cont.)

However, since M is non-singular, so $d_1\|\beta(\alpha)\| \leqslant \|M\beta(\alpha)\| \leqslant d_2\|\beta(\alpha)\|$ for some $d_1, d_2 > 0$, and since the y_k's are l.i., $d_3\|\beta(\alpha)\| \leqslant \left\|\sum_{k=1}^n \beta_k(\alpha) y_k\right\| \leqslant d_4\|\beta(\alpha)\|$, for some $d_3, d_4 > 0$. Therefore, from the equation above, $\|\beta(\alpha)\| = o(\alpha)$, and thus also $\left\|\sum_{k=1}^n \beta_k(\alpha) y_k\right\| = o(\alpha)$.

Along the curve $\alpha \mapsto x_0 + \alpha h + \sum_{k=1}^n \beta_k(\alpha) y_k$, f assumes its local minimum at x_0, so

$$\delta f(x_0; h) = \left.\frac{d}{d\alpha} f\left(x_0 + \alpha h + \sum_{k=1}^n \beta_k(\alpha) y_k\right)\right|_{\alpha=0} = \left.\frac{d}{d\alpha} f\left(x_0 + \alpha h + o(\alpha)\right)\right|_{\alpha=0} = 0.$$

$\square$

Lemma. Let $g_1, \dots, g_n$ be l.i. linear functionals on a vector space X, and let f be a linear functional on X s.t. $f(x) = 0$ for all $x \in X$ s.t. $g_j(x) = 0$ for all $j = 1, \dots, n$. Then, $f \in \mathrm{lin}\{g_1, \dots, g_n\}$.

Proof. Let $G \in \mathscr{L}(X, \mathbb{R}^{n+1})$, where $G_j(x) = g_j(x)$ $(j = 1, \dots, n)$ and $G_{n+1}(x) = f(x)$. Note that $\mathscr{R}(G)$ is a linear subspace of $\mathbb{R}^{n+1}$, and due to the condition on f, it does not intersect $\{(0, \dots, 0, x) \colon x \neq 0\}$, so $\dim \mathscr{R}(G) < n + 1$ and there is a $\lambda \in \mathbb{R}^{n+1} \setminus \{0\}$ s.t. $\lambda_1 g_1(x) + \cdots + \lambda_n g_n(x) + \lambda_{n+1} f(x) = 0$ for all $x \in X$. Since the g_j's are l.i., $\lambda_{n+1} \neq 0$, so dividing by $-\lambda_{n+1}$ gives $f = \tilde{\lambda}_1 g_1 + \cdots + \tilde{\lambda}_n g_n$ for some $\tilde{\lambda}_j$'s. □

From this lemma and Theorem 1, it follows immediately that

Theorem 2 (Lagrange multipliers). Under the conditions of Theorem 1, there exist constants $\lambda_1, \dots, \lambda_n$ s.t.

$$\delta f(x_0; h) + \sum_{i=1}^{n} \lambda_i \delta g_i(x_0; h) = 0 \quad \text{for all } h \in X.$$

Example: Maximum entropy spectral analysis (Burg's) method

Consider the problem of estimating the spectrum Φ of a stationary Gaussian stochastic process, given estimates of the first n autocovariance coefficients. This problem is ill-posed, but one can appeal to the *maximum entropy method* to obtain an estimate:

$$
\begin{array}{cll}
\max\limits_{\Phi \in C[-\pi,\pi]} & H(\Phi) = \ln\sqrt{2\pi e} + \dfrac{1}{4\pi}\displaystyle\int_{-\pi}^{\pi} \ln \Phi(\omega) d\omega & \text{entropy rate of Gaussian process} \\
\text{s.t.} & \dfrac{1}{2\pi}\displaystyle\int_{-\pi}^{\pi} e^{ik\omega}\Phi(\omega) d\omega = c_{|k|}, \quad k = 0, 1, \ldots, n, & \text{autocorrelation coefficients} \\
& \Phi(\omega) \geqslant 0, \quad \text{for all } \omega \in [-\pi, \pi]. & \text{non-negativity constraint}
\end{array}
$$

We will solve this problem using calculus of variations, ignoring the non-negativity constraint (since the solution, as will be seen, is already non-negative). We will assume that the autocorrelation coefficients $c_0, c_1, \ldots, c_n$ are s.t. the problem has feasible solutions.

Example: Maximum entropy spectral analysis (Burg's) method (cont.)
Using Lagrange multipliers, an optimal solution Φ^{opt} should satisfy

$$\frac{1}{4\pi}\int_{-\pi}^{\pi}\frac{1}{\Phi^{\text{opt}}(\omega)}h(\omega)d\omega+\frac{1}{2\pi}\sum_{k=-n}^{n}\lambda_{|k|}\int_{-\pi}^{\pi}e^{ik\omega}h(\omega)d\omega=0,\quad \text{for all } h\in C[-\pi,\pi].$$

Hence, using the fundamental lemma of calculus of variations,

$$\frac{1}{\Phi^{\text{opt}}(\omega)}+2\sum_{k=-n}^{n}\lambda_{|k|}e^{ik\omega}=0 \quad\Leftrightarrow\quad \Phi^{\text{opt}}(\omega)=-\frac{1}{2\sum_{k=-n}^{n}\lambda_{|k|}e^{ik\omega}},$$

where the quantities $\lambda_0,\lambda_1,\dots,\lambda_n$ can be determined from the autocorrelation coefficients $c_0,c_1,\dots,c_n$. This formula shows that the maximum-entropy spectrum corresponds to that of an "auto-regressive process".

Remark. The fact that Φ^{opt} is a maximizer of the optimization problem follows from the concavity of the cost function, and its non-negativity is due to that $H(\Phi)=-\infty$ if Φ is negative inside an interval of $[-\pi,\pi]$ (yielding lower cost than any feasible Φ).