Optimal Control 2018

Kaoru Yamamoto

Optimal Control 2018

- L1: Functional minimization, Calculus of variations (CV) problem
- L2: **Constrained CV problems, From CV to optimal control**
- L3: Maximum principle, Existence of optimal control
- L4: Maximum principle (proof)
- L5: Dynamic programming, Hamilton-Jacobi-Bellman equation
- L6: Linear quadratic regulator
- L7: Numerical methods for optimal control problems

Exercise sessions (20%):
Solve 50% of problems in advance. Hand-in later.
Mini-project (20%):
Study and present your own optimal control problem.
Written take-home exam (60%).

Summary of L1

- $J(y) = \int_a^b L(x, y(x), y'(x))dx,\ y(a) = y_0, y(b) = y_1$.
- First-order necessary condition ⇔ Euler-Lagrange equation
- Alternative form of Euler-Lagrange equation and Hamiltonian
- Weak extrema (necessary conditions are for strong extrema too)
- Variable-endpoint problems

Outline

- Constrained calculus of variations problems
- Second order conditions
- Weierstrass necessary condition for strong extrema
- Cost functional in optimal control problems

Variational problems with constraints

- Euler-Lagrange equation for basic CV problems
  unconstrained except for the boundary conditions
- Equality constraints are now imposed.
  - Integral constraints $\int_a^b M(x, y(x), y'(x))dx = C_0$
  - Non-integral constraints $M(x, y(x), y'(x)) = C_0$

Dido's isopermetric problem

A legend about the foundation of Carthage around 850 B.C.

Dido was allowed to have the land along the North Africa coastline that could be enclosed by an oxhide. She sliced the hide into very thin strips so that she was able to enclose a large area.

- Assume a straight coast line.
- Maximize the area given by
  $J(y) = \int_a^b y(x)dx,\ y : [a, b] \to \mathbb{R}$.
- Constraint:
  - $y(a) = y(b) = 0$,
  - $\int_a^b \sqrt{1 + (y'(x))^2}\, dx = C_0$.

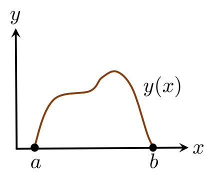

The pendulum

Recall Hamilton's principle of least action in L1. Trajectories of motion for the pendulum are given by solving the following minimization problem:

$$\text{minimize } \int_{t_0}^{t_1} \left(\frac{1}{2}m(\dot{x}^2 + \dot{y}^2) - mgy\right) dt$$
$$\text{subject to } M(x, y) := x^2 + y^2 - l^2 = 0.$$

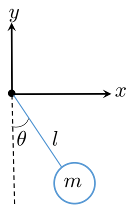

Constrained optimization - Lagrange mulipliers

First-order necessary condition for constrained optimality:

$$\boxed{\nabla f(x^*) + \lambda_1^* \nabla h_1(x^*) + \cdots + \lambda_m^* \nabla h_m(x^*) = 0.}$$

For $x \in \mathbb{R}^2$, (i.e., minimize $f(x_1, x_2)$ subject to $h(x_1, x_2) = 0$)

$$\boxed{\nabla f(x_1^*, x_2^*) = -\lambda \nabla h(x_1^*, x_2^*)}$$

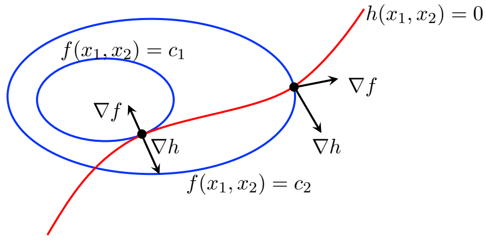

Integral constraints

$$C(y) := \int_a^b M(x, y(x), y'(x))dx = C_0.$$

- Dido's problem, catenary example.
- For η to be admissible, $C(y + \alpha\eta) = C_0$ for $\alpha \approx 0$
- $\Rightarrow \delta C(y, \eta) = 0.$
- $\Rightarrow \int_a^b (M_y(x, y(x), y'(x)) - \frac{d}{dx} M_{y'}(x, y(x), y'(x)))\eta(x)\, dx = 0.$
- $\delta J(y, \eta) = 0$ for every η satisfying the above equation.

$$\int_a^b \left(L_y - \frac{d}{dx} L_{y'}\right) \eta(x)\, dx = 0 \quad \forall \eta \text{ s.t. } \int_a^b \left(M_y - \frac{d}{dx} M_{y'}\right) \eta(x)\, dx = 0.$$

$$\Rightarrow \left(L_y - \frac{d}{dx} L_{y'}\right) + \lambda^* \left(M_y - \frac{d}{dx} M_{y'}\right) = 0 \quad \forall x \in [a, b].$$

Integral constraints cont.

$$\left(L_y - \frac{d}{dx} L_{y'}\right) + \lambda^* \left(M_y - \frac{d}{dx} M_{y'}\right) = 0$$

$$\Longrightarrow (L + \lambda^* M)_y = \frac{d}{dx}(L + \lambda^* M)_{y'}$$

- Euler-Lagrange equation for augmented Lagrangian $L + \lambda^* M$.
- Some gaps in the argument - see [Liberzon 2.5.1].
- We have to be careful with the case y is the extremal of C. (i.e., $M_y - \frac{d}{dx} M_{y'} = 0$)

Example

minimize $J(y) = \int_0^1 L(x)dx$

subject to $C(y) = \int_0^1 \sqrt{1 + (y'(x))^2}\, dx = 1, \quad y(0) = y(1) = 0.$

- The only admissible curve is $y \equiv 0$. $\Rightarrow \quad y^* \equiv 0$ for any L.
- $(M_y - \frac{d}{dx} M_{y'})|_{y=0} = -\frac{d}{dx} \frac{y'(x)}{\sqrt{1+(y'(x))^2}}|_{y=0} = 0.$
- $(L_y - \frac{d}{dx} L_{y'})|_{y=0} = 0$? not necessarily.

Modified augmented cost:

$$\lambda_0^* J + \lambda^* C = \int_a^b (\lambda_0^* L + \lambda^* M)dx, \quad (\lambda_0^*, \lambda^*) \neq (0, 0).$$

- $\lambda_0^* = 0 \Rightarrow y$ is an extremal of C.
- $\lambda_0^* \neq 0 \Rightarrow y$ is an extremal of $J + (\lambda^*/\lambda_0^*)C$.
- λ_0^*: **abnormal multiplier**

Non-integral constraints

$$M(x, y(x), y'(x))dx = 0 \quad \forall x \in [a, b].$$

- Euler-Lagrange eq. for augmented Lagrangian $L + \lambda^*(x)M$.
- Similar to the integral constraint case $(L + \lambda^* M)$
- Instead of the entire interval, the Euler-Lagrange equation holds for every $x \in [a, b]$.
- $\Rightarrow$ A different multiplier for each $x \in [a, b]$.

Second-order conditions

- $J(y + \alpha\eta) = J(y) + \delta J(y, \eta)\alpha + \delta^2 J(y, \eta)\alpha^2 + o(\alpha^2).$
- **Second-order necessary condition** for optimality (**Legendre's condition**)

  $\delta^2 J(y, \eta) \geq 0 \implies L_{y'y'}(x, y(x), y'(x)) \geq 0, \forall x \in [a, b]$
- **Second-order sufficient condition** for optimality

  $\delta J(y) = 0$ and $\delta^2 J(y, \eta) > 0$

  $\implies L_y = \frac{d}{dx} L_{y'}$ and $L_{y'y'}(x, y(x), y'(x)) > 0, \forall x \in [a, b]$

  and $[a, b]$ contains no points conjugate to a.
- Careful arguments required to deal with $o(\alpha^2)$ [Liberzon 2.6].

Legendre's condition and the Hamiltonian maximization

Recall that for the momentum $p := L_{y'}(x, y, y')$ and the Hamiltonian $H(x, y, y', p) := p \cdot y' - L(x, y, y')$,

- $H_{y'} = 0$. $\Rightarrow H$ has a stationary point as a function of y' along an optimal curve $(x, y(x), p$ fixed).
- $H^*(z) := p \cdot z - L(x, y, z)$ then $\frac{dH^*}{dz}(y'(x)) = 0$.
- This stationary point is actually a maximum ($\Rightarrow$ the maximum principle, L3 – L4)

Since $H_{y'y'} = -L_{y'y'} \leq 0$ (Legendre's condition), or

$$\frac{d^2 H^*(z)}{dz^2}(y'(x)) = -L_{y'y'}(x, y(x), y'(x)) \leq 0,$$

if the stationary point is an extremum, it must be a maximum.

Necessary conditions for strong extrema

- Weak minima over $\mathcal{C}^1$ cureves so far
- Stronger notions of local optimality over less regular curves needed
- Strong minima over piecewise $\mathcal{C}^1$ curves
- Continuous y, a finite number of points of discontinous y' – *corner points*
- Such y is a candidate of minima.

Example

Minimize

$$J(y) = \int_{-1}^1 y^2(x)(y'(x) - 1)^2 dx$$

subject to

$$y(-1) = 0, \, y(1) = 1.$$

- Clearly $J(y) \geq 0\, \forall y$.
- We can find $y \in \mathcal{C}^1$ s.t. $J(y) \approx 0$ but not $J(y) = 0$.
- Instead, the curve

$$y(x) = \begin{cases} 0 & \text{if} \quad -1 \leq x < 0 \\ x & \text{if} \quad 0 \leq x < 1 \end{cases}$$

gives $J(y) = 0$.

Necessary conditions for strong extrema

- Weak minima over $\mathcal{C}^1$ cureves so far
- Stronger notions of local optimality over less regular curves needed
- Strong minima over piecewise $\mathcal{C}^1$ curves
- Continuous y, a finite number of points of discontinuous y' – *corner points*
- Such y is a candidate of minima.
- Euler-Lagrange equation (integral form) must hold at all noncorner points. (*extremals, broken extremals*)
- What else?

A perturbation of an extremal with a corner

Corner points

A corner point is a point $c \in [a, b]$ such that $y'(c^-) := \lim_{x \nearrow c} y'(x)$ and $y'(c^+) := \lim_{x \searrow c} y'(x)$ both exists but have different values.

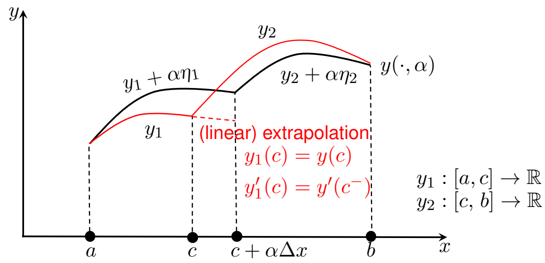

- $\eta_1(a) = \eta_2(b) = 0, \quad \eta_1, \eta_2 \in \mathcal{C}^1$.
- The corner point location not fixed – deviate from c.

Weierstrass-Erdmann corner conditions

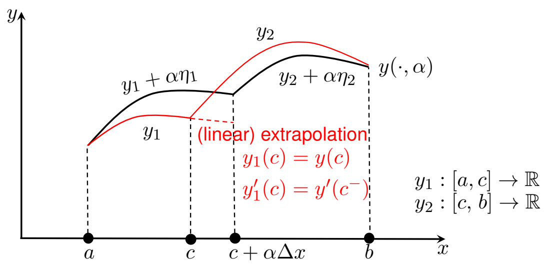

Weierstrass-Erdmann corner conditions

If a curve y is a strong extremum, then $L_{y'}$ and $y'L_{y'} - L$ must be continuous at each corner point of y.

i.e., their discontinuities are *removable.*

Weierstrass excess function

Weierstrass excess function, or **E-function**:

$$E(x, y, z, w) := L(x, y, w) - L(x, y, z) - (w - z)L_z(x, y, z)$$

Weierstrass necessary condition for a strong minimum

$$y \text{ is a strong minimum} \implies E(x, y(x), y'(x), w) \geq 0.$$

for all noncorner points $x \in [a, b]$and all $w \in \mathbb{R}$.

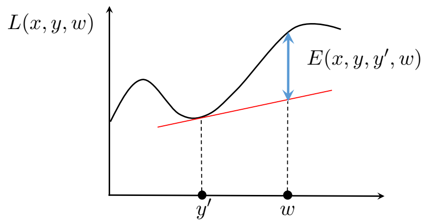

Weierstrass necessary condition and Hamiltonian

$$\begin{aligned} E(x, y, z, w) &= L(x, y, w) - L(x, y, z) - (w - z)L_z(x, y, z) \\ &= zL_z(x, y, z) - L(x, y, z) - (wL_z(x, y, z) - L(x, y, w)) \\ &= H(x, y, z, p) - H(x, y, w, p) \end{aligned}$$

where $p = L_z(x, y, z)$. Hence, the Weierstrass necessary condition implies

$$E(x, y(x), y'(x), w) = H(x, y(x), y'(x), p(x)) - H(x, y(x), w, p(x)) \geq 0.$$

interpretation: if $y(\cdot)$ is an optimal trajectory and $p(\cdot)$ is the corresponding momentum, $\forall x$, $H(x, y(x), \cdot, p(x))$ has a maximum at $y'(x)$.

From calculus of variations to optimal control

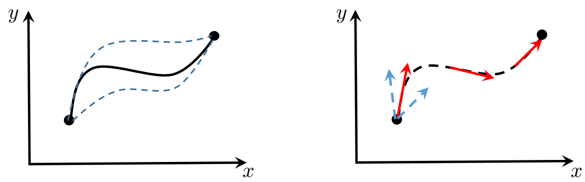

Calculus of variations

- curves given a priori
- curves parameterized by the spacial variable x

Optimal control

- a particle drawing a trace of its motion
- $y' = u$, i.e., optimal control decision at each point
- curves parameterized by time t

Brachistochrone

Find the shortest possible time to travel from one point to the other in a vertical plane.

Calculus of variations ([Liberzon 2.1.4], E1):

$$\text{minimize} \quad J(y) = \int_a^b \frac{\sqrt{1 + (y'(x))^2}}{\sqrt{2gy(x)}} dx$$

$$\text{subject to} \quad y(a) = 0,\ y(b) = y_1.$$

Optimal control:

$$\text{minimize} \quad J(u_1, u_2) = t_1 - t_0 = \int_{t_0}^{t_1} 1 dt$$

$$\begin{aligned} \text{subject to} \quad & (\dot{x}, \dot{y}) = (u_1\sqrt{2gy}, u_2\sqrt{2gy}), \\ & (x(t_0), y(t_0)) = (a, 0), \\ & (x(t_1), y(t_1)) = (b, y_1), \\ & u_1^2 + u_2^2 = 1. \quad \text{(speed constraint)} \end{aligned}$$

Optimal control problem formulation

Find a control $u \in U \subset \mathbb{R}^m$ that minimizes the cost

$$J(u) := \int_{t_0}^{t_f} \underbrace{L(t, x(t), u(t))}_{\text{running cost}} dt + \underbrace{K(t_f, x_f)}_{\text{terminal cost}}$$

subject to

$$\dot{x} = f(t, x, u),\ x(t_0) = x_0,\ x \in \mathbb{R}^n.$$

Cost functional

- Bolza form: $J(u) = \int_{t_0}^{t_f} L(t, x(t), u(t))dt + K(t_f, x_f)$
- Lagrange form: $K \equiv 0$, Mayer form: $L \equiv 0$

Bolza form to Mayer form: introduce an extra state variable x^0 via

$$\dot{x}^0 = L(t, x(t), u(t)), \quad x^0(t_0) = 0.$$

$$\Rightarrow \int_{t_0}^{t_f} L(t, x(t), u(t))dt + K(t_f, x_f) = x^0(t_f) + K(t_f, x_f).$$

Bolza form to Lagrange form:

$$J(u) = \int_{t_0}^{t_f} \left(L(t, x(t), u(t)) + \frac{d}{dt} K(t, x(t)) \right) dt + K(t_0, x_0).$$

$K(t_0, x_0)$: a constant independent of u
$\rightarrow$ can be removed from the optimization problem.

Target set

$$J(u) := \int_{t_0}^{t_f} L(t, x(t), u(t))dt + K(t_f, x_f)$$

- t_0, x_0 are fixed.
- t_f, x_f can be free or fixed, or can belong to some set.
- Captured by introducing a *target set* $S \subset [t_0, \infty) \times \mathbb{R}^n$.
- t_f: the smallest time s.t. $(t_f, x_f) \in S$.

Free-time, fixed-endpoint: $S = [t_0, \infty) \times \{x_1\}, \quad x_1 \in \mathbb{R}^n$.
Fixed-time, free-endpoint: $S = \{t_1\} \times \mathbb{R}^n, \quad t_1 \in [t_0, \infty)$.
Fixed-time, fixed-endpoint: $S = \{t_1\} \times \{x_1\}$.
Free-time, free-endpoint: $S = [t_0, \infty) \times \mathbb{R}^n$.

Calculus of variations vs optimal control

Perturbation

Consider $S = \{t_1\} \times \mathbb{R}^n$, $u \in U = \mathbb{R}^m$ (unconstrained).

$$J(u) = \int_{t_0}^{t_1} L(t, x(t), u(t))dt + K(x(t_1))$$
$$\dot{x} = f(t, x, u), \quad x(t_0) = x_0. \tag{1}$$

- $x = x^* + \alpha\eta$ needs to satisfy (1), but hard to characterize such η.
- u is the design variable – makes more sense to perturb u.
- $u = u^* + \alpha\xi$
- characterize η s.t. the solution of (1) for such u is $x = x^* + \alpha\eta + o(\alpha)$.
- $\dot{\eta} = f_x(t, x^*, u^*)\eta + f_u(t, x^*, u^*)\xi$, $\eta(t_0) = 0$. ([Liberzon 3.4.1])

Augmented cost and Hamiltonian

$$J(u) = \int_{t_0}^{t_1} (L(t, x(t), u(t)) + p(t) \cdot (\dot{x}(t) - f(t, x, u))) dt + K(x(t_1))$$

- $p : \mathcal{C}^1([t_0, t_1] \to \mathbb{R}^n)$.
- Recall the Lagrange multiplier function $\lambda(\cdot)$.
- Also closely related to the momentum defined in L1.

Define the **Hamiltonian** (in optimal control setting) by

$$H(t, x, u, p) := \langle p, f(t, x, u)\rangle - L(t, x, u)$$

$$\Rightarrow J(u) = \int_{t_0}^{t_1} (\langle p, \dot{x}\rangle - H(t, x(t), u(t), p(t)))\, dt + K(x(t_1))$$

We want to compute $\delta J(u^*, \xi)$ of J in this form.

First variation

$$J(u) = \int_{t_0}^{t_1} (\langle p, \dot{x}\rangle - H(t, x(t), u(t), p(t)))\, dt + K(x(t_1))$$

- $J(u^* + \alpha\xi) - J(u^*) = \delta J(u^*, \xi)\alpha + o(\alpha)$
- $\int_{t_0}^{t_1} \langle p(t), \dot{x}(t) - \dot{x}^*(t)\rangle dt$ – integration by parts
- $H(t, x^* + \alpha\eta + o(\alpha), u^* + \alpha\xi, p) - H(t, x^*, u^*, p)$
- $K(x^*(t_1) + \alpha\eta(t_1) + o(\alpha)) - K(x^*(t_1))$

$$\delta J(u^*, \xi) = -\int_{t_0}^{t_1} (\langle \dot{p} + H_x(t, x^*, u^*, p), \eta\rangle + \langle H_u(t, x^*, u^*, p), \xi\rangle)\, dt + \langle K_x(x^*(t_1)) + p(t_1), \eta(t_1)\rangle$$

where $\dot{\eta} = f_x(t, x^*, u^*)\eta + f_u(t, x^*, u^*)\xi, \quad \eta(t_0) = 0$.

First-order necessary condition for optimality

$$\delta J(u^*, \xi) = -\int_{t_0}^{t_1} (\langle \dot{p} + H_x(t, x^*, u^*, p), \eta\rangle + \langle H_u(t, x^*, u^*, p), \xi\rangle)\, dt + \langle K_x(x^*(t_1)) + p(t_1), \eta(t_1)\rangle$$

Pick a special $p = p^*$ s.t.

$$\dot{p}^* = -H_x(t, x^*, u^*, p^*), \quad p^*(t_1) = -K_x(x^*(t_1))$$

$$\Rightarrow \delta J(u^*, \xi) = -\int_{t_0}^{t_1} \langle H_u(t, x^*, u^*, p^*), \xi\rangle dt$$

$\delta J(u^*, \xi) = 0 \; \forall \xi$ implies

$$H_u(t, x^*(t), u^*(t), p^*(t)) = 0 \quad \forall t \in [t_0, t_1].$$

$H(t, x^*(t), \cdot, p^*(t))$ has a stationary point (maximum, in fact) at $u^*(t)$ for all t.

Hamilton's canonical equations

The joint evolution of x^* and p^* is governed by

$$\dot{x}^* = H_p(t, x^*, u^*, p^*)$$
$$\dot{p}^* = -H_x(t, x^*, u^*, p^*)$$

(Note: $H(t, x, u, p) = \langle p, f(t, x, u)\rangle - L(t, x, u)$.)

- p is called the **adjoint vector**.

$$\dot{p}^* = -(f_x(t, x^*, u^*))^T p^* + L(t, x^*, u^*).$$

Compare with $\dot{\eta} = f_x(t, x^*, u^*)\eta + f_u(t, x^*, u^*)\xi$.

- p is also called the **costate** as we can think of p as acting on the state velocity vector by $\langle p, \dot{x}\rangle$.

Necessary conditions for optimality (conjecture)

If $u^*(\cdot)$ an optimal control and $x^*(\cdot)$ the corresponding optimal state trajectory, $\exists p^*$ s.t.:

1) x^* and p^* satisfy, w.r.t. $H(t, x, u, p) = \langle p, f(t, x, u)\rangle - L(t, x, u)$,

$$\dot{x}^* = H_p(t, x^*, u^*, p^*)$$
$$\dot{p}^* = -H_x(t, x^*, u^*, p^*),$$

with $x^*(t_0) = x_0, \quad p^*(t_1) = -K_x(x^*(t_1))$.

2) For each fixed t, the function $u \mapsto H(t, x^*(t), u, p^*(t))$ has a (local) maximum at $u = u^*(t)$:

$$H(t, x^*(t), u^*(t), p^*(t)) \geq H(t, x^*(t), u, p^*(t))$$

for all u near $u^*(t)$ and all $t \in [t_0, t_1]$.