AA203
Optimal and Learning-based Control

Calculus of variations

Stanford University

Autonomous Systems Lab

Roadmap

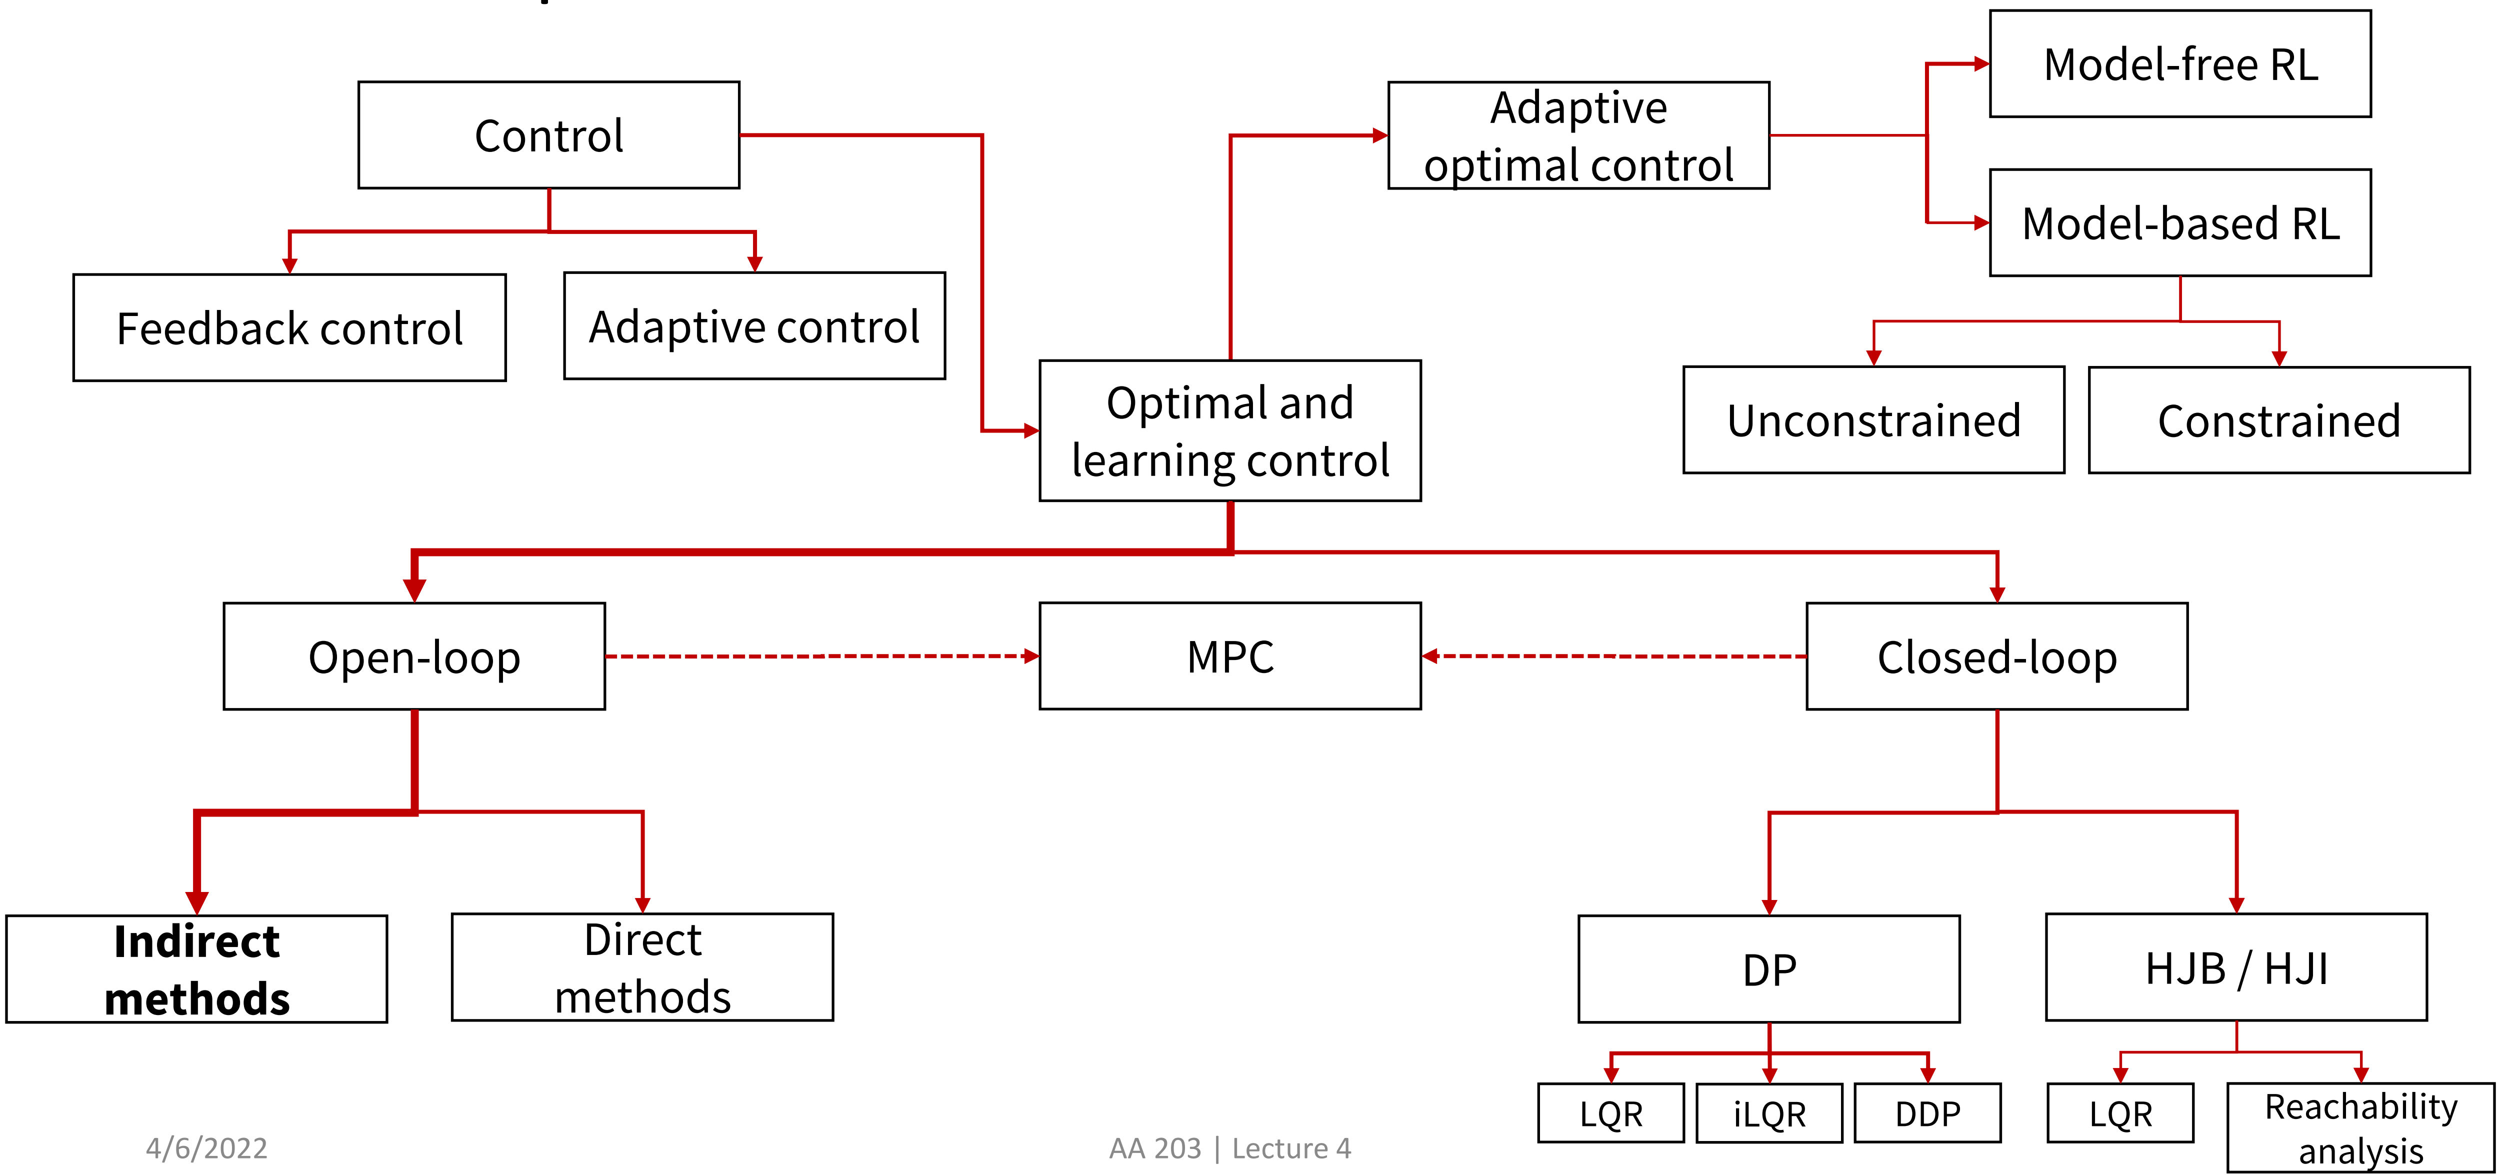

Control
Feedback control
Adaptive control
Optimal and learning control
Adaptive optimal control
Model-free RL
Model-based RL
Unconstrained
Constrained
Open-loop
MPC
Closed-loop
Indirect methods
Direct methods
DP
HJB / HJI
LQR
iLQR
DDP
LQR
Reachability analysis

Indirect methods

Goal: determine necessary conditions for optimality for a general class of optimal control problems

- “Optimize then discretize”
- Sometimes provides more direct (i.e., analytical) path to a solution; otherwise indirect methods enjoy faster convergence with better precision than direct methods (provided you can get them to work…)

Reading:

- D. E. Kirk. *Optimal control theory: an introduction*, 2004.

Key idea

Recall OCP: find an *admissible control* $\mathbf{u}^*$ which causes the system

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t)$$

to follow an *admissible trajectory* $\mathbf{x}^*$ that minimizes the *functional*

$$J = h(\mathbf{x}(t_f), t_f) + \int_{t_0}^{t_f} g(\mathbf{x}(t), \mathbf{u}(t), t)\, dt$$

- For a function, we set gradient to zero to find stationary points, and then investigate higher order derivatives to determine minimum / maximum
- <u>We'll do something very similar for functionals</u>

Calculus of variations (CoV)

- Calculus of variations: generalization of calculus that allows one to find maxima and minima of functionals (i.e., a "function of functions"), by using *variations*

- Agenda:
  1. Introduce new concepts for functionals by appealing to some familiar results from the theory of functions
  2. Apply such concepts to derive the fundamental theorem of CoV
  3. Apply the CoV to optimal control

Preliminaries

- A functional J is a rule of correspondence that assigns to each function $\mathbf{x}$ in a certain class Ω (the “domain”) a real number
  - Example: $J(\mathbf{x}) = \int_{t_0}^{t_f} \mathbf{x}(t)dt$

- J is a linear functional of $\mathbf{x}$ if and only if

$$J\left(\alpha_1\mathbf{x}^{(1)} + \alpha_2\mathbf{x}^{(2)}\right) = \alpha_1 J\left(\mathbf{x}^{(1)}\right) + \alpha_2 J(\mathbf{x}^{(2)})$$

for all $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}$, and $\alpha_1\mathbf{x}^{(1)} + \alpha_2\mathbf{x}^{(2)}$ in Ω
  - Example: previous functional is linear

Preliminaries

To define the notion of (local) maxima and minima, we need a notion of "closeness"

- The norm of a function is a rule of correspondence that assigns to each function $\mathbf{x} \in \Omega$, defined over $t \in [t_0, t_f]$, a real number. The norm of $\mathbf{x}$, denoted by $\|\mathbf{x}\|$, satisfies the following properties:
  1. $\|\mathbf{x}\| \geq 0$, and $\|\mathbf{x}\| = 0$ iff $\mathbf{x}(t) = 0$ for all $t \in [t_0, t_f]$ (positive definite)
  2. $\|\alpha\mathbf{x}\| = |\alpha|\|\mathbf{x}\|$ for all real numbers α (absolute homogeneity)
  3. $\|\mathbf{x}^{(1)} + \mathbf{x}^{(2)}\| \leq \|\mathbf{x}^{(1)}\| + \|\mathbf{x}^{(2)}\|$ (triangle inequality)
- To compare the closeness of two functions $\mathbf{y}$ and $\mathbf{z}$, we let $\mathbf{x}(t) = \mathbf{y}(t)$ - $\mathbf{z}(t)$
  - Example, considering scalar functions $\mathbf{x} \in C^0$:
    $\|\mathbf{x}\|_\infty = \max_{t_0 \leq t \leq t_f} \{|\mathbf{x}(t)|\}$

Extrema for functionals

- A functional J with domain Ω has a local minimum at $\mathbf{x}^*(t) \in \Omega$ if there exists an $\epsilon > 0$ such that
$$J\big(\mathbf{x}(t)\big) \geq J(\mathbf{x}^*(t))$$
for all $\mathbf{x}(t) \in \Omega$ such that
$$\|\mathbf{x}(t) - \mathbf{x}^*(t)\| < \epsilon$$

- Maxima are defined similarly

- To find a minimum we define something similar to the differential of a function

Increments and variations

- The increment of a functional is defined as
$$\Delta J\big(\mathbf{x}(t), \delta\mathbf{x}(t)\big) := J\big(\mathbf{x}(t) + \underbrace{\delta\mathbf{x}(t)}_{\text{Variation of } \mathbf{x}}\big) - J(\mathbf{x}(t))$$

- The increment of a functional can be written as
$$\Delta J(\mathbf{x}, \delta\mathbf{x}) := \delta J(\mathbf{x}, \delta\mathbf{x}) + g(\mathbf{x}, \delta\mathbf{x}) \cdot ||\delta\mathbf{x}||$$
where δJ is *linear* in $\delta\mathbf{x}$. If
$$\lim_{||\delta\mathbf{x}|| \to 0} \{g(\mathbf{x}, \delta\mathbf{x})\} = 0$$
then J is said to be differentiable on $\mathbf{x}$ and δJ is the variation of J at $\mathbf{x}$

The fundamental theorem of CoV

- Let $\mathbf{x}(t)$ be a vector function of t in the class Ω, and $J(\mathbf{x})$ be a differentiable functional of $\mathbf{x}$. *Assume that the functions in Ω are not constrained by any boundaries*. If $\mathbf{x}^*$ is an extremal, the variation of J must vanish at $\mathbf{x}^*$, that is

  $\delta J(\mathbf{x}^*, \delta\mathbf{x}) = 0$ for all admissible $\delta\mathbf{x}$
  (i.e., such that $\mathbf{x} + \delta\mathbf{x} \in \Omega$)

- Proof: by contradiction (see also Kirk, Section 4.1).

Applying CoV

- Let $\mathbf{x}$ be a function in the class of functions with continuous first derivatives. It is desired to find the function $\mathbf{x}^*$ for which the functional

$$J(\mathbf{x}) = \int_{t_0}^{t_f} g(\mathbf{x}(t), \dot{\mathbf{x}}(t), t) dt$$

  has a relative extremum

- Assumptions: $g \in C^2, t_0, t_f$ are fixed, and $\mathbf{x}_0, \mathbf{x}_f$ are fixed

Applying CoV

- Let $\mathbf{x}$ be any element of Ω, and determine the variation δJ from the increment

$$\Delta J(\mathbf{x}, \delta\mathbf{x}) = J(\mathbf{x} + \delta\mathbf{x}) - J(\mathbf{x})$$
$$= \int_{t_0}^{t_f} g(\mathbf{x} + \delta\mathbf{x}, \dot{\mathbf{x}} + \delta\dot{\mathbf{x}}, t)dt - \int_{t_0}^{t_f} g(\mathbf{x}, \dot{\mathbf{x}}, t)dt$$
$$= \int_{t_0}^{t_f} g(\mathbf{x} + \delta\mathbf{x}, \dot{\mathbf{x}} + \delta\dot{\mathbf{x}}, t) - g(\mathbf{x}, \dot{\mathbf{x}}, t)dt$$

- Note that $\dot{\mathbf{x}} = d\ \mathbf{x}(t)/dt$ and $\delta\dot{\mathbf{x}} = d\ \delta\mathbf{x}(t)/dt$

Applying CoV

- Expanding the integrand in a Taylor series, one obtains

$$\Delta J(\mathbf{x}, \delta\mathbf{x}) = \int_{t_0}^{t_f} g(\mathbf{x}, \dot{\mathbf{x}}, t) + \underbrace{\frac{\partial g}{\partial \mathbf{x}}}_{g_{\mathbf{x}}}(\mathbf{x}, \dot{\mathbf{x}}, t)^T \delta\mathbf{x} + \underbrace{\frac{\partial g}{\partial \dot{\mathbf{x}}}}_{g_{\dot{\mathbf{x}}}}(\mathbf{x}, \dot{\mathbf{x}}, t)^T \delta\dot{\mathbf{x}} + o(\delta\mathbf{x}, \delta\dot{\mathbf{x}}) - g(\mathbf{x}, \dot{\mathbf{x}}, t)\, dt$$

- Thus the variation is

$$\delta J = \int_{t_0}^{t_f} g_{\mathbf{x}}(\mathbf{x}, \dot{\mathbf{x}}, t)^T \delta\mathbf{x} + \; g_{\dot{\mathbf{x}}}(\mathbf{x}, \dot{\mathbf{x}}, t)^T \delta\dot{\mathbf{x}} \; dt$$

Applying CoV

- Integrating by parts one obtains

$$\delta J = \int_{t_0}^{t_f} \left[g_{\mathbf{x}}(\mathbf{x}, \dot{\mathbf{x}}, t) - \frac{d}{dt} g_{\dot{\mathbf{x}}}(\mathbf{x}, \dot{\mathbf{x}}, t) \right]^T \delta\mathbf{x}\, dt$$
$$+ \left[g_{\dot{\mathbf{x}}}(\mathbf{x}, \dot{\mathbf{x}}, t)^T \delta\mathbf{x}(t) \right]_{t_0}^{t_f}$$

- Since $\mathbf{x}(t_0)$ and $\mathbf{x}(t_f)$ are given, $\delta\mathbf{x}(t_0) = 0$ and $\delta\mathbf{x}(t_f) = 0$
- If we now consider an extremal curve, applying the CoV theorem yields

$$\delta J = \int_{t_0}^{t_f} \left[g_{\mathbf{x}}(\mathbf{x}^*, \dot{\mathbf{x}}^*, t) - \frac{d}{dt} g_{\dot{\mathbf{x}}}(\mathbf{x}^*, \dot{\mathbf{x}}^*, t) \right]^T \delta\mathbf{x}\, dt = 0$$

For *all* $\delta\mathbf{x}$!

Applying CoV

- Fundamental lemma of CoV: If a function $\mathbf{h}$ is continuous and

$$\int_{t_0}^{t_f} \mathbf{h}(t)^T \delta\mathbf{x}(t)dt = 0$$

for every function $\delta\mathbf{x}$ that is continuous in the interval $[t_0, t_f]$, then $\mathbf{h}$ must be zero everywhere in the interval $[t_0, t_f]$

Applying CoV

- Applying the fundamental lemma, we find that a necessary condition for $\mathbf{x}^*$ to be an extremal is

$$\boxed{g_{\mathbf{x}}(\mathbf{x}^*, \dot{\mathbf{x}}^*, t) \ - \frac{d}{dt} g_{\dot{\mathbf{x}}}(\mathbf{x}^*, \dot{\mathbf{x}}^*, t) = \mathbf{0}}$$

Euler-Lagrange equation

for all $t \in [t_0, t_f]$

- Non-linear, ordinary, time-varying, second-order differential equation with split boundary conditions (at $\mathbf{x}(t_0)$ and $\mathbf{x}(t_f)$)

Example

- Find shortest path between two given points
  - Solution: straight line!

Summary

- A necessary condition for x^* to be an extremal, in the case of *fixed* final time and *fixed* end point, is

$$g_x(x^*, \dot{x}^*, t) - \frac{d}{dt} g_{\dot{x}}(x^*, \dot{x}^*, t) = 0$$

- More generally, for functionals involving several independent functions, a necessary condition for $\mathbf{x}^*$ to be an extremal, in the case of *fixed* final time and *fixed* end points, is

$$g_{\mathbf{x}}(\mathbf{x}^*, \dot{\mathbf{x}}^*, t) - \frac{d}{dt} g_{\dot{\mathbf{x}}}(\mathbf{x}^*, \dot{\mathbf{x}}^*, t) = \mathbf{0}$$

Next class

- More general boundary conditions
- Constrained extrema
- Application to optimal control