AA 203
Optimal and Learning-Based Control
Pontryagin's maximum principle and indirect methods

Spencer M. Richards

Autonomous Systems Laboratory, Stanford University

April 10, 2023
(last updated May 3, 2023)

Agenda

Agenda

1. Geometry and generalizations of first-order NOCs
2. Weak Pontryagin maximum principle in discrete-time
3. Weak Pontryagin maximum principle in continuous-time
4. Pontryagin maximum principle in continuous-time
5. Indirect methods for optimal control
6. Time-optimal control problems

Review: First-order NOCs

$$\begin{aligned} \underset{x\in\mathbb{R}^n}{\text{minimize}} \quad & f(x) \\ \text{subject to} \quad & h(x) = 0 \\ & g(x) \preceq 0 \end{aligned} \qquad L(x,\lambda,\mu) \coloneqq f(x) + \lambda^{\mathsf{T}} h(x) + \mu^{\mathsf{T}} g(x)$$

Theorem (First-order NOCs)

Suppose $x^* \in \mathbb{R}^n$ *is a local minimum of* $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R})$ *subject to* $h(x^*) = 0$ *and* $g(x^*) \preceq 0$ *with* $h \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^m)$ *and* $g \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^r)$. *Moreover, assume*

$$\{\nabla h_i(x^*)\}_{i=1}^m \cup \{\nabla g_j(x^*)\}_{j\in\mathcal{A}_g(x^*)}$$

are linearly independent. Then there exist unique $\lambda^* \in \mathbb{R}^m$ *and* $\mu^* \in \mathbb{R}^r$ *such that*

$$\nabla_x L(x^*,\lambda^*,\mu^*) = 0, \qquad \mu^* \succeq 0, \qquad \mu_j^* = 0,\ \forall j \notin \mathcal{A}_g(x^*),$$

The assumption on the constraint gradients is known as the *linear independence constraint qualification (LICQ)*.

Geometry of first-order NOCs

Tangent cone $\mathcal{T}_{\mathcal{X}}(x)$ "vectors that stay in $\mathcal{X}$"

Normal cone $\mathcal{N}_{\mathcal{X}}(x)$ "vectors that leave $\mathcal{X}$"

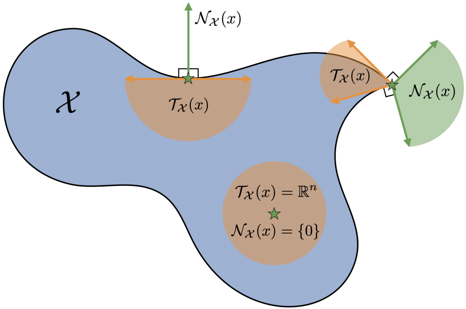

If x^* is a local minimum of f over $\mathcal{X}$, then $-\nabla f(x^*) \in \mathcal{N}_{\mathcal{X}}(x^*)$, i.e., there is no feasible component of $-\nabla f(x^*)$ that would allow us to locally decrease $f(x^*)$.

For convenience, we write "$-\nabla f(x^*) \perp_{x^*} \mathcal{X}$". In other literature, you may see "$-\nabla f(x^*) \perp \mathcal{T}_{\mathcal{X}}(x^*)$".

If $\mathcal{X} = \{x \in \mathbb{R}^n \mid h(x) = 0,\ g(x) \preceq 0\}$ and the LICQ holds at $x^* \in \mathcal{X}$, then

$$\mathcal{T}_{\mathcal{X}}(x^*) = \left\{ d \in \mathbb{R}^n \mid \frac{\partial h}{\partial x}(x^*) d = 0,\ \nabla g_j(x^*)^\mathsf{T} d \leq 0,\ \forall j \in \mathcal{A}_g(x^*) \right\}$$

$$\mathcal{N}_{\mathcal{X}}(x^*) = \left\{ v \in \mathbb{R}^n \mid v = \frac{\partial h}{\partial x}(x^*)^\mathsf{T} \lambda + \frac{\partial g}{\partial x}(x^*)^\mathsf{T} \mu,\ \mu \succeq 0,\ \mu_j = 0, \forall j \notin \mathcal{A}_g(x^*) \right\}$$

Example: A problem with linearly dependent constraints

$$\begin{aligned} \underset{x \in \mathbb{R}^2}{\text{minimize}} \quad & f(x) := x_1 + x_2 \\ \text{subject to} \quad & h_1(x) := (x_1 - 1)^2 + x_2^2 - 1 = 0 \\ & h_2(x) := (x_1 - 2)^2 + x_2^2 - 4 = 0 \end{aligned}$$

At the only feasible point $x^* = 0$, we have

$$\nabla f(x^*) = (1, 1)$$
$$\nabla h_1(x^*) = (-2, 0), \ \nabla h_2(x^*) = (-4, 0)$$

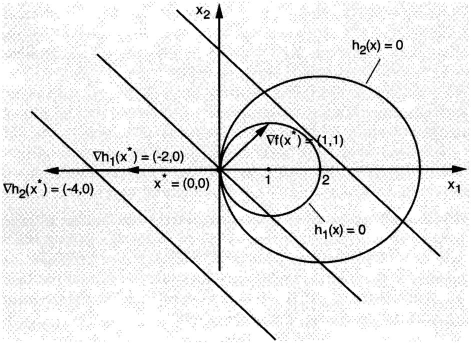

The constraint gradients are linearly dependent (i.e., the LICQ does not hold), so we cannot write $\nabla f(x^*) + \lambda_1^* \nabla h_1(x^*) + \lambda_2^* \nabla h_2(x^*) = 0$.

In essence, the constraints "pinch together" so that just one x^* is feasible, regardless of the objective value.

Fritz John first-order NOCs

Theorem (Fritz John first-order NOCs)

Let $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R})$, $h \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^m)$, *and* $g \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^r)$. *Suppose* $x^* \in \mathbb{R}^n$ *is a local minimum of the problem*

$$\begin{aligned} \underset{x \in \mathcal{S}}{\text{minimize}} \; & f(x) \\ \text{subject to } & h(x) = 0 \; \cdot \\ & g(x) \preceq 0 \end{aligned}$$

Then there exist $(\eta, \lambda^*, \mu^*) \in \{0, 1\} \times \mathbb{R}^m \times \mathbb{R}^r$ *such that*

$$\begin{aligned} (\eta, \lambda^*, \mu^*) &\neq 0 && \textit{non-triviality} \\ -\nabla_x L_\eta(x^*, \lambda^*, \mu^*) &\perp_{x^*} \mathcal{S} && \textit{stationarity} \\ \mu_j^* \geq 0, \;\; \mu_j^* g_j(x^*) &= 0, \;\; \forall j \in \{1, 2, \ldots, r\} && \textit{complementarity} \end{aligned}$$

where $L_\eta(x, \lambda, \mu)$ *is the partial Lagrangian*

$$L_\eta(x, \lambda, \mu) := \eta f(x) + \lambda^\mathsf{T} h(x) + \mu^\mathsf{T} g(x).$$

Fritz John first-order NOCs

Theorem (Fritz John first-order NOCs)

If x^* is a local minimum, there exist $(\eta, \lambda^*, \mu^*) \in \{0,1\} \times \mathbb{R}^m \times \mathbb{R}^r$ such that

$$
\begin{gathered}
(\eta, \lambda^*, \mu^*) \neq 0 \\
-\nabla_x L_\eta(x^*, \lambda^*, \mu^*) \perp_{x^*} \mathcal{S} \\
\mu_j^* \geq 0, \;\; \mu_j^* g_j(x^*) = 0, \;\; \forall j \in \{1, 2, \ldots, r\}
\end{gathered}
$$

where $L_\eta(x, \lambda, \mu)$ is the partial Lagrangian

$$
L_\eta(x, \lambda, \mu) \coloneqq \eta f(x) + \lambda^\mathsf{T} h(x) + \mu^\mathsf{T} g(x).
$$

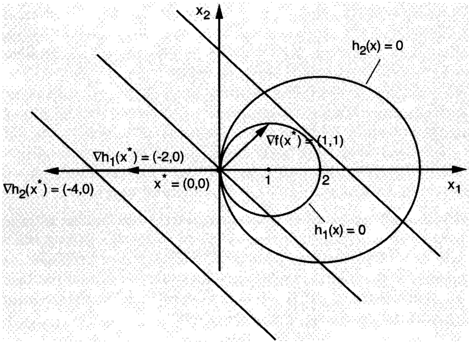

The "abnormal case" $\eta = 0$ yields necessary conditions independent of the objective f.

Corollary

If $\mathcal{S} = \mathbb{R}^n$ and the LICQ holds, then $\eta = 1$ and $\nabla_x L_1(x^*, \lambda^*, \mu^*) = 0$.

Agenda

Course overview

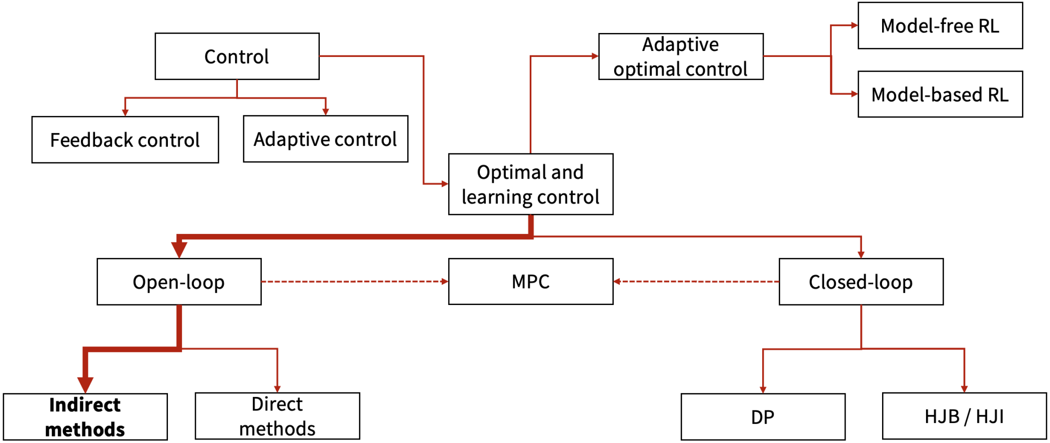

Optimal control problem (discrete-time)

Consider the discrete-time optimal control problem (OCP)

$$
\begin{array}{lll}
\underset{x,u}{\text{minimize}} & \ell_T(x_T) + \displaystyle\sum_{t=0}^{T-1} \ell(t, x_t, u_t) & \text{cost (terminal + stage)} \\
\text{subject to} & x_{t+1} = f(t, x_t, u_t),\ \forall t \in \{0, 1, \ldots, T-1\} & \text{dynamical feasibility} \\
& x_0 = \bar{x}_0 & \text{initial condition} \\
& x_T \in \mathcal{X}_T & \text{terminal condition} \\
& u_t \in \mathcal{U},\ \forall t \in \{0, 1, \ldots, T-1\} & \text{input constraints}
\end{array}
$$

An optimal control $u^* = \{u_t^*\}_{t=0}^{T-1}$ for a specific initial state $\bar{x}_0$ is an *open-loop* input.

An optimal control of the form $u_t^* = \pi^*(t, x_t)$ is a *closed-loop* input.

Lagrangian, Hamiltonian, and the adjoint equation (discrete-time)

The partial Lagrangian is

$$
\begin{aligned}
L_\eta(x,u,p) &= \eta\ell_T(x_T) + \underbrace{p_0^\mathsf{T}(x_0 - \bar{x}_0)}_{\text{initial condition}} + \sum_{t=0}^{T-1}\Big(\eta\ell(t,x_t,u_t) + \underbrace{p_{t+1}^\mathsf{T}(x_{t+1} - f(t,x_t,u_t))}_{\text{dynamical feasibility}}\Big), \\
&= \eta\ell_T(x_T) + p_0^\mathsf{T}(x_0 - \bar{x}_0) + \sum_{t=0}^{T-1}\big(p_{t+1}^\mathsf{T}x_{t+1} - H_\eta(t,x_t,u_t,p_{t+1})\big)
\end{aligned}
$$

with normality $\eta \in \{0,1\}$, Lagrange multipliers $\{p_t\}_{t=0}^T \subset \mathbb{R}^n$, and *Hamiltonian*

$$
H_\eta(t,x,u,p) := p^\mathsf{T} f(t,x,u) - \eta\ell(t,x,u).
$$

Setting $\nabla_{x_t} L(x^*,u^*) = 0$ for $t \in \{0,1,\dots,T-1\}$ yields

$$
p_t^* = \nabla_x H_\eta(t,x_t^*,u_t^*,p_{t+1}^*),\ \forall t \in \{0,1,\dots,T-1\},
$$

which is a *backwards recursion* for the *adjoint* or *co-state* p_t^*.

Transversality and the maximum condition (discrete-time)

The partial Lagrangian is

$$L_\eta(x,u,p) = \eta \ell_T(x_T) + p_0^\mathsf{T}(x_0 - \bar{x}_0) + \sum_{t=0}^{T-1} \big(p_{t+1}^\mathsf{T} x_{t+1} - H_\eta(t, x_t, u_t, p_{t+1})\big)$$

where we left out $x_T \in \mathcal{X}_T$ and $u_t \in \mathcal{U}$. Setting $-\nabla_{x_T} L_\eta(x^*, u^*) \perp_{x_T^*} \mathcal{X}_T$ yields the *transversality condition*

$$-p_T^* - \eta \nabla \ell_T(x_T^*) \perp_{x_T^*} \mathcal{X}_T,$$

and setting $-\nabla_{u_t} L(x^*, u^*) \perp_{u_t^*} \mathcal{U}$ yields the *weak maximum condition*

$$\nabla_u H_\eta(t, x_t^*, u_t^*, p_{t+1}^*) \perp_{u_t^*} \mathcal{U}, \ \forall t \in \{0, 1, \ldots, T-1\}.$$

We refer to this condition as "weak" since it is a necessary, but not sufficient condition for a solution of the problem

$$\operatorname*{maximize}_{u \in \mathcal{U}} H_\eta(t, x_t^*, u, p_{t+1}^*).$$

Pontryagin maximum principle (discrete-time)

Collect these necessary conditions together to get the Pontryagin maximum principle (PMP).

Theorem (Pontryagin maximum principle (discrete-time))

Let (x^*, u^*) be a local minimum of the discrete-time OCP with terminal set $\mathcal{X}_T$ and control set $\mathcal{U}$. Then $\eta \in \{0,1\}$ and $\{p_t^*\}_{t=0}^T \subset \mathbb{R}^n$ exist such that

$$(\eta, p_0^*, p_1^*, \dots, p_T^*) \neq 0 \qquad \textit{non-triviality}$$

$$p_t^* = \nabla_x H_\eta(t, x_t^*, u_t^*, p_{t+1}^*), \ \forall t \in \{0, 1, \dots, T-1\} \qquad \textit{adjoint equation}$$

$$-p_T^* - \eta \nabla \ell_T(x_T^*) \perp_{x_T^*} \mathcal{X}_T \qquad \textit{transversality}$$

$$\nabla_u H_\eta(t, x_t^*, u_t^*, p_{t+1}^*) \perp_{u_t^*} \mathcal{U}, \ \forall t \in \{0, 1, \dots, T-1\} \qquad \textit{maximum condition (weak)}$$

Agenda

1. Geometry and generalizations of first-order NOCs
2. Weak Pontryagin maximum principle in discrete-time
3. Weak Pontryagin maximum principle in continuous-time
4. Pontryagin maximum principle in continuous-time
5. Indirect methods for optimal control
6. Time-optimal control problems

Optimal control problem (continuous-time)

Consider the continuous-time optimal control problem (OCP)

$$
\begin{aligned}
\underset{x,u}{\text{minimize}} \quad & \ell_T(x(T)) + \int_0^T \ell(t, x(t), u(t))\, dt && \text{cost (terminal + stage)} \\
\text{subject to} \quad & \dot{x}(t) = f(t, x(t), u(t)), \ \forall t \in [0, T] && \text{dynamical feasibility} \\
& x(0) = x_0 && \text{initial condition} \\
& x(T) \in \mathcal{X}_T && \text{terminal condition} \\
& u(t) \in \mathcal{U}, \ \forall t \in [0, T] && \text{input constraints}
\end{aligned}
$$

An optimal control $u^*(t)$ for a specific initial state x_0 is an *open-loop* input.

An optimal control of the form $u^*(t) = \pi^*(t, x(t))$ is a *closed-loop* input.

Discretized OCPs

Consider piecewise continuous trajectories such that $x(t) = x(t_k)$ and $u(t) = u(t_k)$ for $t \in [t_k, t_{k+1})$, with $k \in \{0, 1, \ldots, N-1\}$, $t_0 = 0$ and $t_N = T$.

Define $\Delta t_k \coloneqq t_{k+1} - t_k$ such that $\Delta t_k > 0$ for all $k \in \{0, 1, \ldots, N-1\}$.

Consider the discretized OCP

$$
\begin{aligned}
\underset{x,u}{\text{minimize}} \quad & \ell_T(x(t_N)) + \sum_{k=0}^{N-1} \Delta t_k \ell(t_k, x(t_k), u(t_k)) \\
\text{subject to} \quad & x(t_{k+1}) = x(t_k) + \Delta t_k f(t_k, x(t_k), u(t_k)), \ \forall k \in \{0, 1, \ldots, N-1\} \\
& x(t_0) = x_0 \\
& x(t_N) \in \mathcal{X}_T \\
& u(t_k) \in \mathcal{U}, \ \forall k \in \{0, 1, \ldots, N-1\}
\end{aligned}
$$

Discrete-time PMP as a heuristic for continuous-time OCPs

Use the discrete-time PMP on a local minimum (x^*, u^*) of the discretized OCP to get

$$(\eta, p(t_0), p(t_1), \ldots, p(t_N)) \neq 0$$

$$-\frac{(p^*(t_{k+1}) - p^*(t_k))}{\Delta t_k} = \nabla_x H_\eta(t_k, x^*(t_k), u^*(t_k), p^*(t_{k+1})), \ \forall k \in \{0, 1, \ldots, N-1\}$$

$$-p^*(t_N) - \eta \nabla \ell_T(x^*(t_N)) \perp_{x^*(t_N)} \mathcal{X}_T$$

$$\nabla_u H_\eta(t_k, x^*(t_k), u^*(t_k), p^*(t_{k+1})) \perp_{u_t^*} \mathcal{U}, \ \forall k \in \{0, 1, \ldots, N-1\}$$

where we use the *continuous-time Hamiltonian*

$$H_\eta(t, x, u, p) := p^\mathsf{T} f(t, x, u) - \eta \ell(t, x, u).$$

Pontryagin maximum principle (continuous-time, weak)

The above conditions suggest the following continuous-time PMP as $\Delta t_k \to 0$.

Theorem (Pontryagin maximum principle (continuous-time, weak))

Let (x^*, u^*) be a local minimum of the continuous-time optimal control problem with terminal set $\mathcal{X}_T$ and control set $\mathcal{U}$. Then $\eta \in \{0,1\}$ and $p^* : [0,T] \to \mathbb{R}^n$ exist such that

$$(\eta, p(t)) \not\equiv 0 \qquad \textit{non-triviality}$$

$$-\dot{p}^*(t) = \nabla_x H_\eta(t, x^*(t), u^*(t), p^*(t)), \ \forall t \in [0,T] \qquad \textit{adjoint equation}$$

$$-p^*(T) - \eta \nabla \ell_T(x^*(T)) \perp_{x^*(T)} \mathcal{X}_T \qquad \textit{transversality}$$

$$\nabla H_\eta(t, x^*(t), u^*(t), p^*(t)) \perp_{u^*(t)} \mathcal{U}, \ \forall t \in [0,T] \qquad \textit{maximum condition (weak)}$$

"$(\eta, p(t)) \not\equiv 0$" means there exists at least one $t \in [0,T]$ such that $(\eta, p(t)) \neq 0$.

Agenda

1. Geometry and generalizations of first-order NOCs
2. Weak Pontryagin maximum principle in discrete-time
3. Weak Pontryagin maximum principle in continuous-time
4. Pontryagin maximum principle in continuous-time
5. Indirect methods for optimal control
6. Time-optimal control problems

Norms in function spaces

Recall that (x^*, u^*) is a *local minimum* of $J(x^*, u^*)$ if there exists $\varepsilon > 0$ such that $J(x^*, u^*) \leq J(x, u)$ for all (x, u) in the ε-sized norm ball around (x^*, u^*).

In using the discrete-time PMP as a heuristic to obtain the continuous-time PMP, we are implicitly using the $\mathcal{C}^0$*-norm* for both x^* and u^*, i.e.,

$$\|x - x^*\|_{\mathcal{C}^0} := \max_{t \in [0,T]} \|x(t) - x^*(t)\|, \quad \|u - u^*\|_{\mathcal{C}^0} := \max_{t \in [0,T]} \|u(t) - u^*(t)\|.$$

We can strengthen the continuous-time PMP if we use the $\mathcal{C}^0$-norm for x^* and the $\mathcal{L}^1$-norm for u^*, i.e.,

$$\|x - x^*\|_{\mathcal{C}^0} := \max_{t \in [0,T]} \|x(t) - x^*(t)\|, \quad \|u - u^*\|_{\mathcal{L}^1} := \int_0^T \|u(t) - u^*(t)\| \, dt.$$

Strengthening the maximum condition via needle perturbations

In general, the $\mathcal{L}^1$-norm ball for u^* allows for *large pointwise variations* at each time t. Suppose the control set $\mathcal{U}$ is bounded, i.e., $\|u - v\| \leq c$ for all $u, v \in \mathcal{U}$ and some $c > 0$.

Given some $u^* : [0, T] \to \mathcal{U}$, any $\tau \in [0, T)$ and $\varepsilon > 0$ such that $[\tau, \tau + \varepsilon) \subset [0, T]$, and any $v \in \mathcal{U}$, define

$$u(t) = \begin{cases} v, & t \in [\tau, \tau + \varepsilon) \\ u^*(t), & t \in [0, \tau) \cup [\tau + \varepsilon, T] \end{cases}$$

This is a *spatial needle perturbation* of $u^*(t)$. Then it can be shown that

$$\|u - u^*\|_{\mathcal{L}^1} \coloneqq \int_0^T \|u(t) - u^*(t)\|\, dt = \int_\tau^{\tau+\varepsilon} \|v - u^*(t)\|\, dt \leq \int_\tau^{\tau+\varepsilon} c\, dt = \varepsilon c.$$

$$x(T) \approx x^*(T) + \varepsilon d,\ d \in \mathcal{T}_{\mathcal{X}_T}(x^*(T))$$

for small enough ε. Overall, a large spatial perturbation in $u^*(t)$ can correspond to small feasible perturbations to both x^* and u^*.

Pontryagin maximum principle (continuous-time)

The possibility of large spatial control perturbations still corresponding to "feasible neighbours" of (x^*, u^*) suggests the following strengthened PMP.

Theorem (Pontryagin maximum principle (continuous-time))

Let (x^*, u^*) be a local minimum (using the $\mathcal{C}^0$-norm and $\mathcal{L}^1$-norm, respectively) of the continuous-time OCP with terminal set $\mathcal{X}_T$ and bounded control set $\mathcal{U}$. Then $\eta \in \{0,1\}$ and $p^* : [0,T] \to \mathbb{R}^n$ exist such that

$$
\begin{aligned}
(\eta, p^*(t)) &\not\equiv 0 && \textit{non-triviality}\\
-\dot{p}^*(t) &= \nabla_x H_\eta(t, x^*(t), u^*(t), p^*(t)),\ \forall t \in [0,T] && \textit{adjoint equation}\\
-p^*(T) - \eta\, \nabla \ell_T(x^*(T)) &\perp_{x^*(T)} \mathcal{X}_T && \textit{transversality}\\
H_\eta(t, x^*(t), u^*(t), p^*(t)) &= \sup_{u \in \mathcal{U}} H_\eta(t, x^*(t), u, p^*(t)),\ \forall t \in [0,T] && \textit{maximum condition}
\end{aligned}
$$

A rigorous proof relies on variational calculus (Liberzon, 2012; Clarke, 2013).

Example: Minimum fuel for a control-affine system

Consider the continuous-time OCP

$$
\begin{aligned}
\underset{x,u}{\text{minimize}} \quad & \int_0^T \sum_{j=1}^m \alpha_j |u_j(t)| \, dt \\
\text{subject to} \quad & \dot{x}(t) = a(t, x(t)) + \sum_{j=1}^m u_j(t) b_j(t, x(t)), \ \ \forall t \in [0, T] \\
& x(0) = x_0 \\
& x(T) = 0 \\
& -\bar{u} \preceq u(t) \preceq \bar{u}, \ \ \forall t \in [0, T]
\end{aligned}
$$

where $\bar{u} \succ 0$. The Hamiltonian is

$$
H_\eta(t, x, u, p) = p^\mathsf{T} \left(a(t, x) + \sum_{j=1}^m u_j b_j(t, x) \right) - \eta \sum_{j=1}^m \alpha_j |u_j|
$$

Example: Minimum fuel for a control-affine system

The Hamiltonian is

$$H_\eta(t,x,u,p) = a(t,x)^\mathsf{T} p + \sum_{j=1}^{m} \left(u_j b_j(t,x)^\mathsf{T} p - \eta \alpha_j |u_j| \right)$$

The adjoint equation is

$$\dot{p}^* = -\nabla_x H_\eta(t,x^*,u^*,p^*) = -\frac{\partial a}{\partial x}(t,x^*)p^* - \sum_{j=1}^{m} u_j^* \frac{\partial b_j}{\partial x}(t,x^*)p^*$$

The maximum condition is

$$u_j^* = \operatorname*{arg\,max}_{u_j \in [-\bar{u}_j, \bar{u}_j]} \left(u_j b_j(t,x^*)^\mathsf{T} p^* - \eta \alpha_j |u_j| \right) = \begin{cases} -\bar{u}_j, & b_j(t,x^*)^\mathsf{T} p^* < -\eta\alpha_j \\ 0, & b_j(t,x^*)^\mathsf{T} p^* \in [-\eta\alpha_j, \eta\alpha_j] \\ \bar{u}_j, & b_j(t,x^*)^\mathsf{T} p^* > \eta\alpha_j \end{cases},$$

which for $\eta = 1$ is an example of “bang-off-bang” control.

Example: Minimum fuel for a control-affine system

Assume $\eta = 1$, i.e., the "normal" case. Altogether, we have the *boundary value problem (BVP)*

$$\begin{pmatrix} \dot{x}^* \\ \dot{p}^* \end{pmatrix} = \begin{pmatrix} a(t, x^*) + \sum_{j=1}^m u_j^* b_j(t, x^*) \\ -\frac{\partial a}{\partial x}(t, x^*)p^* - \sum_{j=1}^m u_j^* \frac{\partial b_j}{\partial x}(t, x^*)p^* \end{pmatrix}, \quad u_j^* = \begin{cases} -\bar{u}_j, & b_j(t, x^*)^\mathsf{T} p^* < -\alpha_j \\ 0, & b_j(t, x^*)^\mathsf{T} p^* \in [-\alpha_j, \alpha_j] \\ \bar{u}_j, & b_j(t, x^*)^\mathsf{T} p^* > \alpha_j \end{cases},$$

with boundary conditions $x^*(0) = x_0$ and $x^*(T) = 0$.

Transversality did not factor into this problem, since the normal cone of the singleton $\mathcal{X}_T = \{0\}$ is just $\mathbb{R}^n$ (i.e., any direction "leaves" the terminal set).

Agenda

Indirect methods for optimal control

An *indirect method* generally focuses on solving the BVP

$$\begin{pmatrix} \dot{x}^* \\ \dot{p}^* \end{pmatrix} = \begin{pmatrix} f(t, x^*, u^*) \\ -\nabla_x H_\eta(t, x^*, u^*(t, x^*, p^*), p^*) \end{pmatrix}, \quad x^*(0) = x_0, \quad h(x^*(T), p^*(T)) = 0.$$

where $h(x^*(T), p^*(T)) \in \mathbb{R}^n$. The *open-loop* optimal control candidate $u^*(t, x^*(t), p^*(t))$ is then extracted.

The boundary condition $h(x^*(T), p^*(T)) = 0$ is determined by the terminal set constraint $x^*(T) \in \mathcal{X}_T$ and the transversality condition $-p^*(T) - \eta \nabla \ell_T(x^*(T)) \perp_{x^*(T)} \mathcal{X}_T$.

We are implicitly assuming an optimal control exists. Even then, there may be multiple local optima.

Shooting methods

To solve the BVP

$$\begin{pmatrix} \dot{x}^* \\ \dot{p}^* \end{pmatrix} = \begin{pmatrix} f(t, x^*, u^*) \\ -\nabla_x H_\eta(t, x^*, u^*(t, x^*, p^*), p^*) \end{pmatrix}, \quad x^*(0) = x_0, \quad h(x^*(T), p^*(T)) = 0,$$

we consider the associated initial value problem (IVP)

$$\begin{pmatrix} \dot{x}^* \\ \dot{p}^* \end{pmatrix} = \begin{pmatrix} f(t, x^*, u^*) \\ -\nabla_x H_\eta(t, x^*, u^*(t, x^*, p^*), p^*) \end{pmatrix}, \quad x^*(0) = x_0, \quad p^*(0) = p_0.$$

We can integrate the IVP forward in time to get $x^*(T; p_0)$ and $p^*(T; p_0)$, which are parameterized by p_0.

We can use a root-finding method (e.g., bisection search, Newton-Raphson method) to find p_0 such that $h(x^*(T; p_0), p^*(T; p_0)) = 0$. This is called *single shooting* and gives us a solution of the BVP.

Agenda

Time-optimal control problems

Consider the continuous-time OCP

$$
\begin{aligned}
\underset{x,u,T\geq 0}{\text{minimize}}\;\; & \ell_T(T,x(T)) + \int_0^T \ell(t,x(t),u(t))\,dt && \text{cost (terminal + stage)}\\
\text{subject to}\;\; & \dot{x}(t) = f(t,x(t),u(t)),\;\; \forall t\in[0,T] && \text{dynamical feasibility}\\
& x(0) = x_0 && \text{initial condition}\\
& x(T) \in \mathcal{X}_T && \text{terminal condition}\\
& u(t) \in \mathcal{U},\;\; \forall t\in[0,T] && \text{input constraints}
\end{aligned}
$$

The final time T is now a *free variable* (subject to $T \geq 0$).

Time-optimal control problems

Use the change of variables $t(s) = Ts$ with $s \in [0,1]$ to get

$$
\begin{array}{lll}
\underset{(x,t),(u,T)}{\text{minimize}} & \ell_T(t(1), x(1)) + T \int_0^1 \ell(t(s), x(s), u(s))\, ds & \text{cost (terminal + stage)} \\
\text{subject to} & \dot{x}(s) = Tf(t(s), x(s), u(s)),\ \dot{t}(s) = T,\ \forall s \in [0,1] & \text{dynamical feasibility} \\
& x(0) = x_0,\ t(0) = 0 & \text{initial condition} \\
& x(1) \in \mathcal{X}_T & \text{terminal condition} \\
& u(s) \in \mathcal{U},\ T \in [0, \infty),\ \forall s \in [0,1] & \text{input constraints}
\end{array}
$$

To derive a new form of the PMP for time-optimal problems, we apply the fixed final time PMP to the problem above, where we treat t and T as a new state and input, respectively.

Deriving the time-optimal PMP

Applying the fixed final time PMP gives us the Hamiltonian

$$\tilde{H}_\eta(s, x, t, u, T, p, \lambda) = T(H(t, x, u, p) + \lambda),$$

where $H(t, x, u, p)$ is the usual Hamiltonian, and λ is the adjoint for the new "state" $t(s) = Ts$. Taking derivatives with respect to (x, t) yields the adjoint equations

$$\frac{dp^*}{ds} = -T^* \nabla_x H(t, x^*, u^*, p^*), \quad \frac{d\lambda^*}{ds} = -T^* \frac{\partial H}{\partial t}(t, x^*, u^*, p^*),$$

which by the chain rule with $\frac{dt}{ds} = T$ become

$$\dot{p}^* = -\nabla_x H(t, x^*, u^*, p^*), \quad \dot{\lambda}^* = -\frac{\partial H}{\partial t}(t, x^*, u^*, p^*).$$

Since t has no terminal constraint, we have the transversality conditions

$$-p^*(1) - \eta \nabla_x \ell_T(t(1), x^*(1)) \perp_{x^*(1)} \mathcal{X}_T, \quad -\lambda^*(1) - \eta \nabla_T \ell_T(t(1), x^*(1)) = 0.$$

which after using $t = sT$ gives us

$$-p^*(T) - \eta \nabla_x \ell_T(T^*, x^*(T)) \perp_{x^*(T)} \mathcal{X}_T, \quad -\lambda^*(T^*) = \eta \nabla_T \ell_T(T^*, x^*(T)).$$

Deriving the time-optimal PMP

Applying the fixed final time PMP gives us the Hamiltonian

$$\tilde{H}_\eta(s, x, t, u, T, p, \lambda) = T(H(t, x, u, p) + \lambda),$$

where $H(t, x, u, p)$ is the usual Hamiltonian, and λ is the adjoint for the new "state" $t(s) = Ts$

We are considering the absolute value norm for T, and $[0, \infty)$ is unbounded. So we use the maximum condition for u^* and the weak maximum condition for T^* to get

$$\nabla_T \tilde{H}_\eta(t, x^*, u^*, p^*) \perp_{T^*} [0, \infty) \implies H(t, x^*, u^*, p^*) + \lambda^* = 0,$$

where we have assumed $T^* > 0$ to get that the normal cone is just $\{0\}$. Evaluating this condition at $t = T^*$ gives us

$$H(T^*, x^*(T^*), u^*(T^*), p^*(T^*)) = -\lambda^*(T^*) = \eta \nabla_t \ell_T(T^*, x^*(T)),$$

which is the additional boundary condition we need for free final time T^*.

Time-optimal PMP

Collecting all of the conditions we derived above gives us the free final time PMP.

Theorem (Pontryagin maximum principle (continuous-time, free final time))

Let (x^*, u^*, T^*) be a local minimum (using the $\mathcal{C}^0$-norm, $\mathcal{L}^1$-norm, and absolute value, respectively) of the continuous-time OCP with terminal set $\mathcal{X}_T$, bounded control set $\mathcal{U}$, and free final time $T \geq 0$. Then $\eta \in \{0, 1\}$ and $p^* : [0, T^*] \to \mathbb{R}^n$ exist such that

$$(\eta, p^*(t)) \not\equiv 0$$ *non-triviality*

$$-\dot{p}^*(t) = \nabla_x H_\eta(t, x^*(t), u^*(t), p^*(t)), \ \forall t \in [0, T^*]$$ *adjoint equation*

$$-p^*(T^*) - \eta \nabla \ell_T(T^*, x^*(T^*)) \perp_{x^*(T)} \mathcal{X}_T$$ *transversality*

$$H_\eta(t, x^*(t), u^*(t), p^*(t)) = \sup_{u \in \mathcal{U}} H_\eta(t, x^*(t), u, p^*(t)), \ \forall t \in [0, T^*]$$ *maximum condition*

$$H_\eta(T^*, x^*(T^*), u^*(T^*), p^*(T^*)) = \eta \frac{\partial \ell_T}{\partial T}(T^*, x^*(T^*))$$ *maximum condition (boundary)*

Next class

Direct methods for optimal control
(i.e., solving discretized optimal control problems directly)

References

F. Clarke. *Functional Analysis, Calculus of Variations and Optimal Control*. Springer, 2013.

D. Liberzon. *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press, 2012.