Reading group: Calculus of Variations and Optimal Control Theory by Daniel Liberzon

Alexandre Vieira

29th May 2017

Optimal Control Problem

$$\min_{u \in U} J(u) := \int_{t_0}^{t_f} L(t, x(t), u(t))dt + K(t_f, x_f)$$
$$\text{s.t. } \dot{x}(t) = f(t, x(t), u(t))$$
$$(t_0, x(t_0)) \in S_0$$
$$(t_f, x(t_f)) \in S$$

- $x(\cdot) \in \mathbb{R}^n$: the state
- u : the control
- $U \subseteq \mathbb{R}^m$: the control set
- L : the running cost
- K : the terminal cost
- $(t_0, x(t_0))$: the initial time and state
- $(t_f, x(t_f))$: the final time and state
- S_0 : the initial set
- S : the target set

Variational approach

Let us apply some results of the Calculus of variation on the previous problem, where:

- $U = \mathbb{R}^m$
- $S_0 = \{t_0\} \times \{x_0\}$
- $S = \{t_f\} \times \mathbb{R}^m$
- $K = K(x_f)$

Exactly as we did before, one uses perturbation of the optimal solution to find necessary conditions. But this time, the perturbation will be on the optimal control.

Variational approach: linearization

Let u^* be an optimal control : $J(u^*) \leq J(u)$ for all piecewise continuous controls u. Consider:

$$u = u^* + \alpha\xi$$

where ξ is a piecewise continuous function for $[t_0, t_f]$ to $\mathbb{R}^m$ and α a real parameter. This gives rise to a perturbed state:

$$x(t, \alpha) = x^*(t) + \alpha\eta(t) + o(\alpha)$$

where, obviously, $\eta(t_0) = 0$. Deriving it according to α:

$$x_\alpha(t, 0) = \eta(t), \ \forall t \in [t_0, t_f]$$

We differentiate this according to time:

Variational approach: linearization

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

We rewrite more compactly as:

$$\dot{\eta} = A_*(t)\eta + B_*(t)\xi, \;\; \eta(t_0) = 0$$

where

$$A_*(t) = f_x|_*(t) = f_x(t, x^*(t), u^*(t))$$
$$B_*(t) = f_u|_*(t) = f_u(t, x^*(t), u^*(t))$$

Remark

This is the linearization of the original system around the optimal trajectory; cf. Sontag's book

Variational approach: augmented cost

Now, how to deal with the *equality constraint* $\dot{x} = f(t, x, u)$? Via augmented cost!

$$J(u) = \int_{t_0}^{t_f} \left[L(t, x(t), u(t)) + \langle p(t), \dot{x}(t) - f(t, x(t), u(t))\rangle\right] dt + K(x_f)$$

for some $\mathscr{C}^1$ function p to be selected later. Once again, we introduce the Hamiltonian:

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

such that the augmented cost becomes:

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

In order to find necessary condition, we need to compute the first variation $\delta J|_{u^*}$.

Variational approach: first variation

Notation

From now on, we use $\approx$ to write approximation up to order 1.

Let us recall:

$$J(u) - J(u^*) = J(u^* + \alpha\xi) - J(u^*) \approx \delta J|_{u^*}(\xi)\alpha$$

We write first order approximations of the three components in J:

$$K(x(t_f)) - K(x^*(t_f)) = K(x^*(t_f) + \alpha\eta(t_f) + o(\alpha)) - K(x^*(t_f)) \approx \langle K_x(x^*(t_f)), \alpha\eta(t_f)\rangle$$

$$H(t, x, p, u) - H(t, x^*, u^*, p) \approx \langle H_x(t, x^*, u^*, p), \alpha\eta\rangle + \langle H_u(t, x^*, u^*, p), \alpha\xi\rangle$$

$$\int_{t_0}^{t_f} \langle p(t), \dot{x}(t) - \dot{x}^*(t)\rangle dt \approx \langle p(t_f), \alpha\eta(t_f)\rangle - \int_{t_0}^{t_f} \langle \dot{p}(t), \alpha\eta(t)\rangle dt$$

Variational approach: first variation

Combining all of this, we have:

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

where η is related to ξ through the linearization found earlier:

$$\dot{\eta} = A_*(t)\eta + B_*(t)\xi, \ \eta(t_0) = 0$$

Now, the first order condition says that we must have, for all ξ, $\delta J|_{u^*}(\xi) = 0$. But we haven't made any choice concerning p so far!

Variational approach: first variation

Let p^* be the solution of the differential equation:

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

thus, we are left with:

$$\delta J|_{u^*}(\xi) = -\int_{t_0}^{t_f} \langle H_u(t, x^*, u^*, p^*), \xi \rangle dt = 0$$

true for all ξ. This, in turn, implies that:

$$\forall t \in [t_0, t_f], \;\; H_u(t, x^*(t), u^*(t), p^*(t)) = 0$$

Variational approach: first variation

Summary

Let u^* be the optimal solution to the optimal control problem and x^* the associated state. Then there exists a function p^* (called the *adjoint state*) such that:

$$\dot{x} = H_p|_*$$

$$\dot{p} = -H_x|_*$$

$$p(t_f) = K(x^*(t_f))$$

$$H_u|_* = 0$$

Variational approach: critique

However, this method has several drawbacks:

- We never took care of the control set U (since here, it is $\mathbb{R}^m$): it may be a problem when constructing perturbation of u^*.
- Target set: the perturbation never took into consideration the fact that we must reach a certain prescribed target set
- The perturbation were taken here *small* (α was thought as small!). We would like to consider also broader perturbations.

Pontryagin's Maximum Principle

$$\begin{aligned} \min_{u(\cdot)\in U} \; & J(u) := \int_{t_0}^{t_f} L(t, x(t), u(t))dt + K(t_f, x_f) \\ \text{s.t. } & \dot{x}(t) = f(t, x(t), u(t)), \\ & x(t_0) \in S_0, \\ & x(t_f) \in S \end{aligned}$$

Theorem for fixed initial time

Let $u^* : [t_0, t_f] \to U$ be an optimal control and let $x^* : [t_0, t_f] \to \mathbb{R}^n$ be the corresponding optimal state trajectory. Then there exist a function $p^* : [t_0, t_f] \to \mathbb{R}^n$ and a constant $p_0^* \leq 0$ satisfying $(p_0^*, p^*(t)) \neq 0$ for all $t \in [t_0, t_f]$ and having the following properties:

Pontryagin's Maximum Principle

Reading group

Alexandre Vieira

Variational approach

Pontryagin

Existence of optimal control

Exercises

Theorem for fixed initial time

1. x^* and p^* satisfy the equations:

$$\dot{x}^* = H_p(t, x^*, u^*, p^*, p_0^*)$$

$$\dot{p}^* = -H_x(t, x^*, u^*, p^*, p_0^*)$$

where the Hamiltonian $H : \mathbb{R} \times \mathbb{R}^n \times U \times \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ is defined as:

$$H(t, x, u, p, p_0) = \langle p, f(t, x, u) \rangle + p_0 L(t, x, u)$$

2. $x(t_0) \in S_0$, $x(t_f) \in S$,

$$p(t_0) \perp T_{x^*(t_0)} S_0 \text{ and } p(t_f) - p_0^* \frac{\partial K}{\partial x}(t_f, x(t_f)) \perp T_{x^*(t_f)} S$$

Theorem for fixed initial time

3. $H(t, x^*(t), u^*(t), p^*(t), p_0^*) \geq H(t, x^*(t), u(t), p^*(t), p_0^*)$ for all $u(t) \in U$ and $t \in [t_0, t_f]$
4. $\frac{d}{dt} H(t, x^*(t), u^*(t), p^*(t), p_0^*) = \frac{\partial}{\partial t} H(t, x^*(t), u^*(t), p^*(t), p_0^*)$

Remark : Assume $S = \{x \in \mathbb{R}^n : h_1(x) = ... = h_{n-k}(x) = 0\}$ (a k codimensional manifold), where all h_i are smooth. Then, $p \perp T_x S$ actually means:

$$\langle p, d \rangle = 0, \ \forall d \in T_x S$$

where

$$T_x S = \{d \in \mathbb{R}^n : \langle \nabla h_i(x), d \rangle = 0, \ i = 1, ..., n-k\}$$

Example: Double integrator

We apply the Maximum Principle to the time-optimal control problem (i.e. $L \equiv 1$, $K \equiv 0$) of the system:

$$\ddot{x} = u, \ u(t) \in [-1, 1]$$

that we represent by the state-space equations:

$$\dot{x}_1 = x_2, \ \dot{x}_2 = u$$

with $x_1(t_0)$ and $x_2(t_0)$ are known. The Hamiltonian is $H = p_1 x_2 + 2u + p_0$. According to the Maximum Principle, the costate p^* must satisfy the adjoint equation:

$$\begin{pmatrix} \dot{p}_1^* \\ \dot{p}_2^* \end{pmatrix} = \begin{pmatrix} 0 \\ -p_1^* \end{pmatrix}$$

Thus, there exists constants c_1 and c_2 such that $p_2^*(t) = -c_1 t + c_2$, so p_2 is a linear function of time.

Example: Double integrator

Next, from the Hamiltonian maximisation condition and the fact that $U = [-1, 1]$, we have:

$$u^*(t) = \text{sign}\,(p_2^*(t)) = \begin{cases} 1 & \text{if } p_2^*(t) > 0 \\ -1 & \text{if } p_2^*(t) < 0 \\ ? & \text{if } p_2^*(t) = 0 \end{cases}$$

Since p_2 is a linear function of time (and we can prove it is not identically 0), it crosses 0 at most once, so u will switch between values -1 and 1: this is what we call the *bang-bang property*.

Example: Linear Systems

We can derive the same expression for more general linear systems:

$$\dot{x} = Ax + Bu$$

with $U = [-1, 1]^m$. If we denote by b_i the columns of B, we prove in the same way as before that:

$$u_i(t) = \operatorname{sign}(\langle p(t), b_i \rangle)$$

Thus, the function $B^\intercal p(t)$ will tell us what value in $\{-1, 1\}^m$ $u^*(t)$ will take on $[t_0, t_f]$. $B^\intercal p(t)$ is called the *switching function*.

Does the optimal control exist?

So far, we have only necessary conditions to find optimal *candidates*. But are we even sure an optimal solution exists?
The next theorem addresses this problem (and this is not an easy one...).

Does the optimal control exist?

Theorem

Suppose that U is compact and that S is accessible from S_0 (i.e., there exists a control leading from S_0 to S). Let $\mathcal{U}$ be the set of controls with value in U joining S_0 and S. We also suppose that:

1. there exists a positive scalar b such that the trajectory x_u associated to $u \in \mathcal{U}$ is uniformly bounded by b on $[t_0, t_f]$, as long with t_f. It means:

$$\exists b > 0; \forall u \in \mathcal{U}, \forall t \in [t_0, t_f], \; t_f + \|x_u(t)\| \leq b$$

2. For all $(t, x) \in \mathbb{R}^{1+n}$, the set $V(t, x) = \left\{ \begin{pmatrix} f(t, x, u) \\ L(t, x, u) + \gamma \end{pmatrix} | u \in U, \gamma \geq 0 \right\}$ is convex

So there exists an optimal control u on $[t_0, t_f]$ such that the corresponding trajectory joins S_0 and S in time t_f with minimal cost.

Exercises

- Use the Maximum Principle to derive necessary conditions for the unconstrained linear-quadratic optimal problem:

$$\min_{u(\cdot)\in\mathbb{R}^m} \int_0^{t_f} \left(x(t)^\intercal Q(t)x(t) + u(t)^\intercal W(t)u(t)\right) dt$$
$$\text{s.t. } \dot{x}(t) = A(t)x(t) + B(t)u(t) + r(t)$$
$$x(0) = x_0$$

where $Q \in \mathbb{R}^{n\times n}$ and $W \in \mathbb{R}^{m\times m}$ are symmetric positive semi-definite matrices. (Answer: $u(t) = W(t)^{-1}B(t)p(t)$, $\dot{p} = A^\intercal p + Qx$, $p(T) = -Qx(T)$.)

Exercises

- Consider the linear system $\dot{x}(t) = A(t)x(t) + B(t)u(t) + r(t)$, $x(0) = x_0$. The problem here is called the tracking problem: we want a solution $x(\cdot) \in \mathbb{R}^n$ of the previous system tracking on $[0, T]$ a given $\mathscr{C}^1$ trajectory $\xi(\cdot) \in \mathbb{R}^n$, starting for a point ξ_0.
  We introduce the error $z(t) = x(t) - \xi(t)$. We want to minimize the following quadratic cost:

$$J(u) = z(T)^\intercal Qz(t) + \int_0^T (z(t)^\intercal Qz(t) + u(t)^\intercal Wu(t))\, dt$$

1. Write this problem as an optimal control problem on z (differentiate z to obtain the corresponding ODE).
2. Use the Maximum Principle to obtain the necessary conditions of optimality (they are here also sufficient)
3. Application: use this for the oscillator $\ddot{x} + x = u$, $x(0) = 0$, $\dot{x}(0) = 1$, to follow the curve $(\cos(t), \sin(t))$ on $[0, 2\pi]$. Do a numerical implementation.