Lecture 3. The maximum principle

In the last lecture, we learned calculus of variation (CoV). The key idea of CoV for the minimization problem

$$\min_{u \in U} J(u)$$

can be summarized as follows. 1) Assume u^* is a minimizer, and choose a one-parameter variation u_ϵ s.t. $u_0 = u_*$ and $u_\epsilon \in U$ for ϵ small. 2) The function $\epsilon \mapsto J(u_\epsilon)$ has a minimizer at $\epsilon = 0$. Thus it satisfies the first and second order necessary conditions

$$\frac{\partial J}{\partial \epsilon}|_{\epsilon=0} = 0, \quad \frac{\partial^2 J}{\partial \epsilon^2}|_{\epsilon=0} \geq 0$$

because $\epsilon \mapsto J(u_\epsilon)$ is only a scalar valued function on real line. To facilitate the computation, we introduced the δ operator which satisfies the following properties:

P1) δ is a differential operator: it satisfies the chain rule, composition rule, etc.

P2) in our setting, δ commutes with the integration and differentiation operators, i.e., $\delta \int = \int \delta$ and $\delta \dot{x} = \frac{d}{dt}\delta x$.

P3) If a function $u \mapsto J(u)$ has a minimum at u_*, then the first variation vanishes $\delta J(u_*) = 0$ and the second variation is non-negative $\delta^2 J(u_*) \geq 0$.

If you're not very satisfied with the δ notation, then you can always use the naive approach. But trust me, you'll get exactly the same results.

Then we derived the necessary conditions of the Lagrangian problem:

$$\min \int_0^T L(q, \dot{q})\mathrm{d}t$$
$$q(0) = q_0, \quad q(T) = q_1$$

Namely,

$$\frac{\partial L}{\partial q} - \frac{d}{dt}\frac{\partial L}{\partial \dot{q}} = 0.$$
$$L_{\dot{q}\dot{q}} \geq 0$$

Using this equation, we were able to solve a number of interesting questions.

Then, for the Lagrangian problem, we introduced an important coordinate transform, i.e., the canonical transform

$$p = \frac{\partial L}{\partial \dot{q}}$$

together with the introduction of a function called the Hamiltonian

$$H(q, p) = p^\top \dot{q} - L(q, \dot{q})$$

and showed that in the new coordinate system (p, q), Lagragian equation has a much simpler form, that is, the canonical/Hamiltonian equation.

Finally we tried to use CoV to solve the optimal control

$$\min \varphi(x(T)) + \int_0^T L(x, u)\mathrm{d}t$$

subject to

$$\dot{x} = f(x, u), \quad x(0) = x_0.$$

Remember that after introducing

$$H(x, u, p) = p^\top f(x, u) - L(x, u)$$

we were at the step

$$\delta J = -\int_0^T (\dot{p}^\top + H_x)\delta x + H_u \delta u \mathrm{d}t + (\varphi_x + p(T)^\top)\delta x(T) \tag{1}$$

Then by requiring p to satisfy the ODE

$$\dot{p} = -H_x^\top(x_*, u_*)$$

with boundary condition

$$p(T) = -\varphi_x^\top(x_*(T))$$

the first variation becomes

$$\delta J(u_*) = -\int_0^T H_u(x_*, u_*, p)\delta u \mathrm{d}t$$

now that δu is arbitrary, since we suppose we have no constraints, then

$$H_u(x_*(t), u_*(t), p(t)) = 0, \quad \forall t \in [0, T].$$

Let's now derive the second order necessary condition. For that, we shall use (1):

$$\begin{aligned}\delta^2 J = &-\int_0^T \delta x^\top H_{xx}\delta x + \delta u^\top H_{xu}\delta x + (\dot{p}^\top + H_x)\delta^2 x + \delta x^\top H_{ux}\delta u + \delta u^\top H_{uu}\delta u + H_u \delta^2 u \mathrm{d}t \\ &+ (\varphi_x + p(T)^\top)\delta^2 x(T) + \delta x(T)^\top \varphi_{xx}(x(T))\delta x(T)\end{aligned}$$

Note that there won't be any terms involving δp or $\delta \dot{p}$ since δp is multiplied by a zero term – remember that p is added through

$$p^\top(t)(\delta \dot{x} - f_x \delta x - f_u \delta u) = 0.$$

We choose the variation such that $\delta^2 u = 0$ (e.g., $u = u_* + \epsilon v$), then at the optimal point,

$$\begin{aligned}\delta^2 J &= \delta x(T)^\top \varphi_{xx}(x(T))\delta x(T) + \int_0^T -\delta x^\top H_{xx}\delta x - \delta u^\top H_{xu}\delta x - \delta x^\top H_{ux}\delta u - \delta u^\top H_{uu}\delta u \mathrm{d}t \\ &= \delta x(T)^\top \varphi_{xx}(x(T))\delta x(T) + (\varphi_x + p(T)^\top)\delta^2 x(T) - \int_0^T \begin{bmatrix}\delta x \\ \delta u\end{bmatrix}^\top \begin{bmatrix}H_{xx} & H_{xu} \\ H_{ux} & H_{uu}\end{bmatrix}\begin{bmatrix}\delta x \\ \delta u\end{bmatrix} \mathrm{d}t\end{aligned}$$

In order that $\delta^2 J \geq 0$, we must have (the Hessian of $H(\cdot, u)$):

$$H_{uu}(x_*(t), u_*(t), p_*(t)) \leq 0, \quad \forall t \in [0, T].$$

because wild oscillation in δu may cause negligible oscillation in δx. Thus H_{uu} dominates $\delta^2 J$.

Putting together the first and second order conditions

$$\begin{aligned}H_u(x_*(t), u_*(t), p(t)) &= 0, \quad \forall t \in [0, T]. \\ H_{uu}(x_*(t), u_*(t), p_*(t)) &\leq 0, \quad \forall t \in [0, T]\end{aligned}$$

we may guess that

$$u_*(t) \in \mathrm{argmax}_u H(x_*(t), p(t), u).$$

This formula is quite interesting, which says that although the input space (functional space) is infinite dimensional, the optimal $u_*(t)$ is can be obtained through a finite dimensional optimization problem! Based on the previous deduction, we make the following conjecture.

Conjecture 1. *Assume that for the system $\dot{x} = f(x, u)$, f and L are C^2 in both x and u, and that u is unconstrained. Assume $x(0)$ is fixed and $x(T)$ is free. Then there exists some $p(\cdot)$ and that on the optimal solution, we have*

$$\begin{aligned}\dot{x}_* &= H_p^\top(x_*, u_*) \\ \dot{p} &= -H_x(x_*, u_*)\end{aligned}$$

with boundary conditions

$$x_*(0) = x_0, \quad p(T) = -\varphi_x^\top$$

and

$$u_*(t) \in \mathrm{argmax}_u H(x_*(t), u), \quad \forall t \in [0, T].$$

Remark 1. In this conjecture, we need twice continuous differentiability because while taking second order variation, we have to compute the second order partial derivatives. This is restrictive and seems unnecessary – in the formulas, we used only first order partial derivatives of f, L, φ w.r.t. x. So a good question is, does this also hold for f, L, φ which are only piecewise C^1 w.r.t. x? Another drawback is that the conjecture requires u to be constraint free. In practice, this is not very useful because we almost always have constraints on the input. e.g., $|u| \leq 1$. Notice that for this kind of input, when u is at the boundary, say $u = 1$, the variation near this u cannot be arbitrary – it can not exceeds 1. Thus the above reasoning is no longer valid.

However, this conjecture is very close to the final form of the maximum principle which we now introduce.

Theorem 1. *Consider the system $\dot{x} = f(x, u)$ with cost function*

$$J(u) = \varphi(x(t_f)) + \int_0^{t_f} L(x, u)\mathrm{d}t$$

and boundary constraint

$$x(t_f) \in M \subseteq \mathbb{R}^n$$

Assume f, φ and L are C^1 in x. Let $(x^*(\cdot), u^*(\cdot))$ correspond to the optimal solution to the minimization problem

$$\min_{u \in \mathcal{U}_{\mathrm{ad}}} J(u)$$

in which $\mathcal{U}_{\mathrm{ad}} = \{u : [0, t_f] \to U \subseteq \mathbb{R}^m\}$. Define the Hamiltonian function $H(x, u, p, p_0) = p^\top f(x, u) + p_0 L(x, u)$. Then there exists a function $p^* : [0, t_f] \to \mathbb{R}^n$ and a constant $p_0^* \in \{0, -1\}$, satisfying $(p_0^*, p^*(t)) \not\equiv (0, 0)$ such that

1) $(x^*(\cdot), p^*(\cdot))$ satisfies the canonical equation

$$\begin{aligned} \dot{x} &= H_p^\top \\ \dot{p} &= -H_x^\top \end{aligned}$$

with initial condition $x^*(0) = x_0$. The second equation is called the costate equation, and p is the costate.

2) The transversality condition holds:

$$p^*(t_f) + \varphi_x^\top(x^*(t_f)) \perp T_{x^*(t_f)} M.$$

3) The maximum principle holds:

$$H(x^*(t), u^*(t), p^*(t), p_0^*) = \max_{u \in U \subseteq \mathbb{R}^m} H(x^*(t), u, p^*(t), p_0^*) = \text{constant} \tag{2}$$

for all $t \in [0, t_f]$. In particular, this constant is zero if t_f is free.

This theorem was conjectured by the Pontrayagen group in the late 1950s. But it took quite some time before the theorem was finally rigorously proved. Before going on, let's stop for a while and appreciate a bit the theorem.

Discussions

- There's no smoothness requirement of f and L on u – we only require C^1 of f, φ, L on x. This is a big improvement compared to our conjecture. And there's no PDE (as we will see later, there is another method for optimal control called dynamic programming, which requires PDE). This means that the numerical computational load should be quite manageable.

- Compared to the conjecture, there is a new constant p_0 in the Hamiltonian which is either 0 or -1. When $p_0 = 0$, we call the minimizer u_* an abnormal extremal. Abnormal extremals are bad because they are hard to compute. This is because in the Hamiltonian H, the information of the Lagrangian is lost. As a result, the maximum principle tells you only limited information. You might expect some criteria which can help you exclude the existence of abnormal extremals.

Unfortunately, as far as I know, such criteria do not exist for general systems. The good news is, in most control systems, abnormal extremals do not exist and checking the non-existence of abnormal extremals are usually not difficult. A simple method for doing this is to first setting $p_0 = 0$ and then derive a contradiction by showing that $p(t) \equiv 0$.

- Transversality conditions of the costate equation. When the terminal state is free, $M = \mathbb{R}^n$, then the transversality condition is simply

$$p^*(t_f) = -\varphi_x^\top(x(t_f)).$$

With this, we can solve the costate equation backward. When the terminal state is fixed, however, M is a singleton, and the transversality condition does not tell us any information about the terminal condition of $p^*(t_f)$. Instead, one has to use other information to solve p. Typically, one has to solve a boundary value problem resorting to the fact that the initial and terminal states of the system are fixed. In general, boundary value problem is harder to solve than a Cauchy problem. In practice, it's common to relax the boundary hard constraint to a "soft constraint". For example, for a steering problem where the desired terminal state is a fixed point x_1, we consider a cost

$$\tilde{J}(u) := k|x(T) - x_1|^2 + \int_0^T L(x, u)\mathrm{d}t.$$

By adjusting (e.g., increasing) the parameter k, you argue that the terminal state $x(T)$ can be made close to the desired terminal state.

- The maximum principle is a finite dimensional optimization problem. As we said, this is quite remarkable since this turns an infinite minimization problem to a finite dimensional one. For low dimensional problems, sometimes you can compute explicitly from the maximum principle the optimal control $u^*(t)$ as a function of $x^*(t)$ and $p^*(t)$. Then you substitute $u^*(x^*, p^*)$ back to the canonical equation and solve x^* and p^* and at the same time you also get u^*. It is worth mentioning that the maximum principle is a maximization while the optimal control problem is a minimization. This is nothing mysterious as I said, you can swap the sign of H and then you get a minimization principle. Another thing to mention is that the maximum principle says that the Hamiltonian is constant along the optimal trajectory. But you should be careful that this constant is not the optimal cost of the system. For general nonlinear systems, normally, in order to get the optimal cost, you have to substitute the optimal solution to the cost and integrate. For unconstrained linear systems (very limited case!), the optimal cost is easy to find. It's also worth mentioning that when the terminal time t_f is free, this constant is 0 which is a quite interesting result.

Examples

Before we prove this theorem, let's see how should we use the maximum principle. In maximum principle, the first step is to

1) Determine the type of the optimal control problem. e.g., is the terminal state fixed? is the terminal time free? And also we need to verify that the assumptions of the MP are met, e.g., smoothness.

2) Write the Hamiltonian, the costate equation, and the transversality condition – if there's any.

3) Next, solve the maximization and differential equations, which can be done either concurrently or separately, numerically or analytically.

We illustrate this first by a toy example.

Example 1. Consider

$$\begin{aligned}\dot{x}_1 &= -x_1 + u\\ \dot{x}_2 &= x_1\end{aligned}$$

with initial condition $x = (1, 0)$ and cost

$$J = x_2(1).$$

The constraint on u is $|u| \leq 1$.

To solve this, first we define the Hamiltonian $H(x, u, p, p_0) = p_1(-x_1 + u) + p_2 x_1$, which does not depend on p_0. Then we write the costate equation

$$\dot{p}_1 = -H_{x_1} = p_1 - p_2$$
$$\dot{p}_2 = -H_{x_2} = 0$$

with

$$\begin{bmatrix} p_1(1) \\ p_2(1) \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \perp \mathbb{R}^2$$

thus

$$p_1(1) = 0, \quad p_2(1) = -1.$$

By maximum principle, the optimal law should be taken as

$$u_*(t) = \text{sign} p_1(t).$$

But we are not done yet, we have to compute $p_1(t)$. From the costate equation and transversality condition, we get immediately

$$p_2(t) \equiv -1$$

and

$$p_1(t) = e^t p_1(0) + \int_0^t e^{t-\tau} \mathrm{d}\tau = e^t p_1(0) + e^t - 1$$

Using the boundary condition, we get

$$e p_1(0) + e - 1 = 0 \Rightarrow p_1(0) = \frac{1-e}{e}$$

which results in

$$p_1(t) = e^{t-1} - 1.$$

The optimal controller is drawn as in Figure

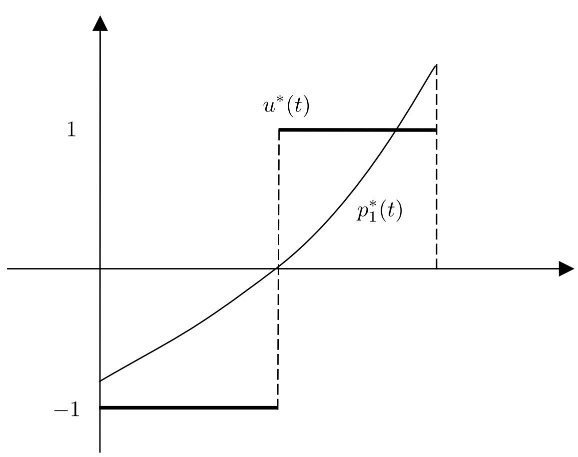

Figure 1: Optimal controller has a switching

The next problem is somehow less artificial.

Minimum fuel control

Suppose that we are to land a lunar rover to the moon. The dynamics of this model is described by

$$\ddot{y} = -g + u$$

where y is the height of the lander, $g \geq 0$ the gravitational acceleration, and u the trust, which can be up or down and is bounded $|u| \leq 1$, and $0 < g < 1$. Note that here we assume the mass of the

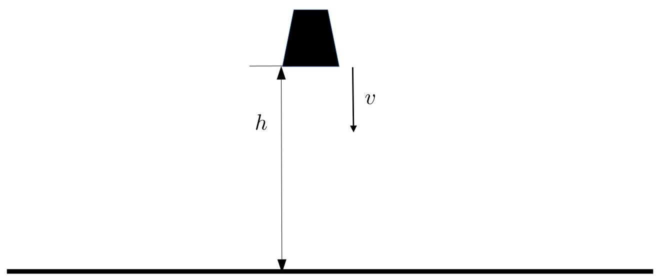

Figure 2: Lunar lander

lander is 1 (fuel loss is neglected). The initial height of the lander is $y(0) = h > 0$ and initial velocity $\dot{y}(0) = v < 0$. In order that the problem is feasible, assume h is large, otherwise the lander can never land with zero velocity.

The objective is to find an optimal control law which minimizes the fuel consumption

$$J = \int_0^{t_f} |u| \mathrm{d}t$$

with t_f free, and which drives the system to the final state $y(t_f) = \dot{y}(t_f) = 0$.

Rewrite the system model as

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -g + u \end{aligned}$$

with initial and terminal conditions $(x_1(0), x_2(0) = (h, v)$, $(x_1(t_f), x_2(t_f)) = (0, 0)$. The Hamiltonian is $H(x, u, p, p_0) = p_1 x_2 + p_2(-g + u) + p_0|u|$, and the costate equation

$$\begin{aligned} \dot{p}_1 &= 0 \\ \dot{p}_2 &= -p_1 \end{aligned}$$

Then $p_1(t) = c_1$ and $p_2(t) = -c_1 t + c_2$ for some constants c_1 and c_2. Since the terminal state is fixed, for the moment we don't know the terminal condition of the costate equation.

Since p_0 appears in the Hamiltonian in a nontrivial manner, we need to exclude abnormal extremals first. If $p_0 = 0$, then

$$u^*(t) = \text{sign}(p_2(t))$$

and using the maximum principle, since t_f is free, we have for all t,

$$H(x^*(t), u^*(t), p^*(t), p_0^*) = c_1 x_2^*(t) + (-c_1 t + c_2)(-g + u^*(t)) = 0. \tag{3}$$

In particular, at $t = t_f$,

$$(-c_1 t_f + c_2)(-g + 1) = 0$$

or $-c_1 t_f + c_2 = 0$. Thus there is no switching and $u^* \equiv 1$ (at the final stage, $u^*(t_f)$ must be non-negative). This only happens when the height and velocity satisfies

$$1 - g = \frac{v^2}{2h}.$$

In this situation, $u^* \equiv 1$ is the only controller possible, which is of course optimal. So there is an abnormal extremal.

Now let $p_0 = -1$. In this case, if we maximize $\max_u H(x, u, p) = p_1 x_2 + p_2(-g + u) - |u|$, we will get

$$u^*(t) = \begin{cases} -1, & p_2 < -1 \\ 0, & -1 \le p_2 < 1 \\ 1, & p_2 \ge 1 \end{cases}$$

As before, u must be positive near t_f, i.e., it must be in the phase $p_2 = -c_1 t + c_2 \geq 1$, for all t near t_f. For this to be true, if $c_1 \geq 0$, then $p_2(t) \geq 1$ for all t and there is no switching and $u_* = 1$. As we mentioned before, this happens only when $1 - g = \frac{v^2}{2h}$. Let $c_1 < 0$. There might be two switchings, when $p_2(t)$ crosses -1 or 1. Since when $c_2 \geq 1$ there's no switching, consider $c_2 \leq 1$. In this case, there is at least one switching when

$$-c_1 t_2 + c_2 = 1$$

or $t_2 = \frac{c_2-1}{c_1}$. In order to have two switchings, there should exist some t_1 satisfying $-c_1 t_1 + c_2 = -1$, or $t_1 = \frac{1+c_2}{c_1}$, which requires $c_2 < -1$ and u^* should be taken as

$$u^*(t) = \begin{cases} -1, & 0 \leq t \leq t_1 \\ 0, & t_1 < t \leq t_2 \\ 1, & t_2 < t \leq t_f \end{cases}.$$

Using $x_2(t_f) = 0$, we can obtain the equality

$$v + (-g-1)t_1 - g(t_2 - t_1) + (1-g)(t_f - t_2) = 0,$$

From this and $H(x^*(t), u^*(t), p^*(t)) \equiv 0$, we can solve for $c_2 = \frac{c_1 v - 1}{1+g} > -1$ since $v < 0$ as assumed, a contradiction. Thus $-1 \leq c_2 < 1$, and there is only one switching at t_2. The corresponding optimal control is

$$u^*(t) = \begin{cases} 0, & 0 < t \leq t_2 \\ 1, & t_2 < t \leq t_f \end{cases}.$$

To find t_2, use the terminal condition $x_1(t_f) = x_2(t_f) = 0$:

$$\begin{aligned} v - gt_2 + (1-g)(t_f - t_2) &= 0 \\ vt_2 - \frac{1}{2}gt_2^2 + \frac{1}{2}(1-g)(t_f - t_2)^2 &= h \end{aligned}$$

from which we find solve for t_2, t_f and then c_1, c_2. To conclude, the lander should first follow a free fall until v and h satisfies $1 - g = \frac{v^2}{2h}$, and then use its full power to propel the lander. See Figure 3.

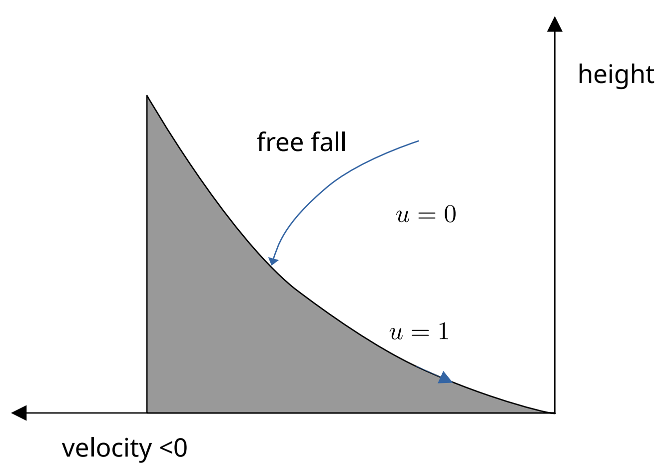

Figure 3: Moon lander