Calculus of Variations, Lecture 7, 3/20/2017

These notes discuss:

A) an alternative perspective on optimal control, via the Pontryagin Maximum Principle

B) sensitivity analysis for pde-constrained optimization, using the adjoint eqn

C) we'll return to the Pontryagin Max Prin, showing that it can be obtained by sensitivity analysis of our optimal control problem

---

A) The Pontryagin Max Principle (PMP) provides an alternative perspective (different from the HJ eqn) on optimal control

A good reference: L Hocking, Optimal Control: An Introduction to the Theory with Applications, Oxford Univ Press 1991. (A really lovely book, with lots of examples.)

Some advantages of PMP:

a) In high dimensions it may be much more efficient numerically than solving the HJ eqn

b) It sometimes leads to explicit (or nearly explicit) solns

c) It provides a very concrete interpretation for the characteristics of the HJ eqn

d) It provides a framework for proving existence of optimal controls (and for introducing "relaxed" formulations where necessary) — though we won't have time for this aspect. (One source for this: RV Gamkrelidze, Principles of Optimal Control Theory, Plenum Press 1978.)

To see the main idea in a simple, transparent setting, consider the "min arrival time problem"

$$u(x) = \min\{\text{time to arrive at } \partial D, \text{ starting from } x \in D \text{ and travelling with speed} \le 1\}$$

for which the HJ eqn is $|\nabla u| = 1$ in D, $u = 0$ at ∂D; solution formula is

$$u(x) = \min_{z \in \partial D} \text{dist}(x, z).$$

There are good alg-s for solving the HJ eqn; basically the build the level-sets one by one starting at $u=0$ (see eg J Sethian's book Level Set Methods & Fast Marching Methods, Cambridge Univ Press)

But using the PMP in this setting amounts to using the solution formula, ie using that optimal choices involves traveling at const velocity in a straight line; the reduces task to search for $z \in \partial D$ st dist (x, z) is smallest. For $D \subset \mathbb{R}^n$ this is an $n-1$ dim'l opt-n, rather than soln of a pde in $D \subset \mathbb{R}^n$.

I'll derive the PMP formally, by a min/max argument (we'll get a more honest derivation later via sensitivity analysis). I'll focus on the finite-time horizon problem

$$u(x,t) = \max_{\alpha \in A} \int_t^T h(y(s), \alpha(s))\, ds + g(y(T))$$

with state eqn $\quad y_s = f(y(s), \alpha(s))$ for $t < s < T$, $y(t) = x$.

To start, we write it as a max/min, using a new vector-valued unknown $\pi(s)$ (taking values in $\mathbb{R}^n$ if $\vec{y}(t) \in \mathbb{R}^n$)

$$u(x,t) = \max_{\substack{\alpha \\ y(t)=x}} \min_{\pi(s)} \left\{ \int_t^T \pi(s)\cdot[f(y(s),\alpha(s)) - \tfrac{dy}{ds}]\,ds + \int_t^T h(y(s),\alpha(s))\,ds + g(y(T)) \right\}$$

where now $y(t)=x$, $\alpha(s)\in A$ are the only constraints in the maximization (the min over $\pi(s)$ enforces the state eqn). Assuming max min = min max (my arg't is "formal" because of this assumption) we get, after integn by parts,

$$u(x,t) = \min_{\pi(s)} \max_{\substack{y(t)=x \\ \alpha(s)\in A}} \left\{ \int_t^T y\cdot\frac{d\pi}{ds} + f(y,\alpha)\cdot\pi + h(y,\alpha)\,ds + g(y(T)) + x\cdot\pi(t) - y(T)\cdot\pi(T) \right\}$$

Now max over $\alpha(s)\in A$: if $\pi(s)$ + $y(s)$ are fixed, best $\alpha(s)$ maximizes $f\cdot\pi + h$. So

$$u(x,t) = \min_{\pi(s)} \max_{y(t)=x} \left\{ \int_t^T [y\cdot\frac{d\pi}{ds} + H(\pi,y)]\,ds + g(y(T)) + x\cdot\pi(t) - y(T)\cdot\pi(T) \right\}$$

with

$$(*) \qquad H(p,x) = \max_{\alpha \in A} \{ p \cdot f(x,\alpha) + h(x,\alpha) \}$$

(note: our HJ eqn, obtained via dynamic programming, was $u_t + H(\nabla u, x) = 0$).

Observe the following property of H:

$$(1) \qquad \frac{\partial H}{\partial p} = f(x, \alpha_0)$$ where α_0 is optimal for the defn (*) of H.

(Proof: let $\alpha_0 = \alpha_0(p,x)$ be the optimal α as fn of p and x. Then

$$H(p,x) = p \cdot f(x, \alpha_0(p,x)) + h(x, \alpha_0(p,x))$$

$$\Rightarrow \quad \frac{\partial H}{\partial p} = f(x, \alpha_0(p,x)) + \left(p \cdot \frac{\partial f}{\partial \alpha} + \frac{\partial h}{\partial \alpha} \right) \Bigg|_{\alpha = \alpha_0(p,x)} \frac{\partial \alpha_0}{\partial p}$$

$$= f(x, \alpha_0(p,x))$$

since the 2nd term vanishes (because α_0 is optimal.) So: if we choose control at time s to be $\alpha_0(\pi(s), y(s))$ then state eqn $\frac{dy}{ds} = f$ becomes

$$(2) \qquad \frac{dy}{ds} = \nabla_p H(\pi(s), y(s))$$

Returning now to the min/max expression for x, let's evaluate the inner max (wrt $y(s)$). Since the only constraint is $y(t) = x$, $y(s)$ is arbitrary for $s > t$, so $y \cdot \frac{d\pi}{ds} + H(\pi, y)$ must be maximized in y for each s; this gives

$$(3) \qquad \frac{d\pi}{ds} = -\frac{\partial H}{\partial y}$$

Arguing similarly to optimize over $y(T)$ we get

$$(4) \qquad \pi(T) = \nabla g\,(y(T))$$

Collecting (1)–(4) we get the PMP for this class of problems:

- $y(s)$ and $\pi(s)$ solve

$$\frac{dy}{ds} = \nabla_\pi H(\pi(s), y(s))$$

$$\frac{d\pi}{ds} = -\nabla_y H(\pi(s), y(s))$$

for $t < s < T$

- initial cond'n for $y(s)$ is known: $y(t) = x$; but initial cond'n for $\pi(s)$ is <u>not</u> known; rather, it is determined implicitly (or at least constrained) by the final-time cond'n

$\pi(T) = \nabla g(y(T))$.

- at each s, the control $\alpha(s)$ should be the value α_0 st $\pi \cdot f(y, \alpha_0) + h(y, \alpha_0)$ is max (with $\pi = \pi(s)$, $y = y(s)$).

Since H is being evaluated at $(\pi(s), y(s))$ and $u_t + H(\nabla u, x) = 0$, it's natural to guess that $\nabla u(y(s), s) = \pi(s)$ along the solution. This is true; the proof rests on the fact that the ODE's for $y(s)$ + $\pi(s)$ are in fact the <u>characteristic equations</u> for our HJ eqn. (Exercise: prove that $\nabla u(y(s), s) = \pi(s)$.)

Can the PMP be used to solve problems? Yes, but it's usually nontrivial since we must somehow get an initial condn for $\pi(s)$. (Typically this might be done numerically, by some sort of search or "shooting" scheme.)

Here's a simple example:

$$\min_{\alpha(s)} \int_0^T \tfrac{1}{2}\alpha^2(s)\, ds + \tfrac{1}{2} y^2(T)$$

with state eqn $y_s = y + \alpha$, $y(0) = x$ (here

$x, y, \alpha \in \mathbb{R}$). Evidently: we prefer $y(T)$ to be small, which requires nonzero α (since $\alpha = 0 \Rightarrow$ exponential growth), but we also prefer $\int \alpha^2$ to be small. (This problem involves minimization not maximization; the PMP is as derived above except that in defining H we must <u>min</u> over α.)

Evidently, in this setting

$$g(y) = \tfrac{1}{2}y^2; \quad f(y,\alpha) = y + \alpha; \quad h(x,\alpha) = \tfrac{1}{2}\alpha^2$$

$$H(p,x) = \min_\alpha \left\{ p(x+\alpha) + \tfrac{1}{2}\alpha^2 \right\} = px - \tfrac{1}{2}p^2$$

and the PMP gives

$$\frac{dy}{ds} = y(s) - \pi(s) \qquad y(0) = x$$

$$\frac{d\pi}{ds} = -\pi(s) \qquad \pi(T) = y(T)$$

It's clear we should be able to find the soln by taking $\pi(0) = \pi_0$ as an unknown, solving for $\pi(s)$ and $y(s)$, then asking which π_0 gives $\pi(T) = y(T)$. (Exercise: do this.)

B) Sensitivity analysis, for pde-constrained optimization.

In multivariable calculus we learn to solve constrained optimizations by the method of Lagrange multipliers. In optimal control our primary constraint is the state eqn and $\pi(s)$ is something like a Lagrange multiplier for it. We'll show in segment (C) how to make sense of this.

But first: let's build intuition + technique by discussing sensitivity analysis for PDE-constrained problems (where the sensitivity is wrt some coefft of the pde). We'll focus on the following model problem: given $D \subset \mathbb{R}^n$, choose $a(x)$ to solve

$$\min_{a(x)} \int_D |u_a - g|^2 \, dx$$

where $g(x)$ is given, and u_a solves elliptic PDE

$$(**) \qquad -\Delta u + b \cdot \nabla u + a(x) u = 1 \text{ in } D$$
$$u = 0 \text{ at } \partial D$$

(Here b is a fixed vector field. It could be zero, but when it is nonzero the operator $\mathcal{L} = -\Delta u + b\cdot\nabla u$ is not self-adjoint.)

My main goal is to ask: what is the gradient (wrt to coefft $a(x)$) of the objective

$$F(a) = \int_D |u_a - g|^2\,dx \qquad ?$$

Explain what this means: suppose we consider a 1-par family $a(x,t) = a_0(x) + t\,\dot a(x)$; assume pde (**) is uniquely solvable (this is a hypothesis on a_0!) so assoc soln $u_a = u_0 + t\dot u + \cdots$. Finding the ($L^2$) gradient of F means finding a function $\xi(x)$ st

$$\frac{d}{dt}\Big|_{t=0} F(a_0 + t\dot a) = \int_D \xi(x)\,\dot a(x)\,dx$$

$$= \langle \xi, \dot a \rangle_{L^2(D)}$$

The utility of this is clear:

- if a_0 is optimal for $\min_a \int_D |u_a - g|^2$ then ξ must be 0
- if $\xi \neq 0$ then using $\dot a = -\xi$ we are

assured that $t \to F(a_0 + t\dot{a})$ is a decreasing fn of t at $t=0$ (ie this gives a descent direction).

Claim: For the functional F defined above, the L^2 gradient ξ at $a = a_0(x)$ is

$$\xi = -u_0(x)\, p_0(x)$$

where $p_0(x)$ solves the "adjoint eqn"

$$-\Delta p_0 - \operatorname{div}(b\, p_0) + a_0\, p_0 = 2(u_0 - g) \text{ in } D$$
$$p_0 = 0 \text{ at } \partial D$$

Proof: we must show that

$$\frac{d}{dt}\Big|_{t=0} F[a_0 + t\dot{a}] = \int_D \xi(x)\, \dot{a}(x)\, dx$$

for any $\dot{a}(x)$. Differentiating the defn of F we get

$$\frac{d}{dt}\Big|_{t=0} F[a_0 + t\dot{a}] = \int_D 2(u_0 - g)\dot{u}\, dx$$

Differentiating the pde we get

$$-\Delta\dot{u} + b\cdot\nabla\dot{u} + a_0\,\dot{u} = -\dot{a}\, u_0 \quad \text{in } D$$
$$\dot{u} = 0 \quad \text{at } \partial D$$

Using the defn of p_0 we get

$$\int_D z(u_0-g)\,\dot{u}\,dx = \int_D [-\Delta p_0 - \mathrm{div}(b\,p_0) + a_0 p_0]\,\dot{u}$$

$$= \int_D (-\Delta\dot{u} + b\cdot\nabla\dot{u} + a_0\dot{u}]\,p_0\,dx$$

(integrating by parts, ie using that LHS of p_0 eqn is adjt of LHS of $\dot{u}$ eqn)

$$= -\int_D \dot{a}\,u_0\,p_0\,dx$$

$$= \int_D \xi(x)\,\dot{a}(x)\,dx$$

when $\xi = -u_0 p_0$.

Our model problem was of course just an example. But the same basic idea has been used in lots of settings; for a recent example see eg M Zahr, P-O Persson, J Wilkening, "A fully discrete adjoint method for optimization of flow problems on deforming domains with time-periodicity constraints," Computers + Fluids 139 (2016) 130-147 [or arxiv:1512.00616]

(C) Returning to the Pontr Max Prin, let's derive it using <u>sensitivity analysis</u>.

We'll use the idea of segment (B), though the constraint is now an ODE (the state eqn) rather than a pde.

We work (as <u>we did earlier</u>) with the example

$$u(x,t) = \max_{\alpha(s)} \int_t^T h(y(s), \alpha(s))\, ds + g(y(T))$$

but for simplicity I'll assume now that $\alpha(s)$ is unrestricted (rather than, say, $\alpha(s) \in A$ for some set A). The state eqn is, as usual,

$$y_s = f(y(s), \alpha(s)) \text{ for } t < s < T, \quad y(t) = x.$$

We want to take the 1st variation of the objective; in other words, given $\alpha_0(s)$ and the assoc soln $y_0(s)$ of the state eqn, we want to consider perturbed controls $\alpha(s, \varepsilon) = \alpha_0(s) + \varepsilon \dot{\alpha}(s)$ and the assoc perturbed solns of state eqn

$$y(s,\varepsilon) = y_0(s) + \varepsilon \dot{y}(s) + \mathcal{O}(\varepsilon^2)$$

and we want to find $\xi(s)$ s.t.

$$\left.\frac{d}{d\varepsilon}\right|_{\varepsilon=0} \int_t^T h(y(s,\varepsilon), \alpha(s,\varepsilon))\,ds + g(y(T,\varepsilon))$$
$$= \int_t^T \xi(s)\,\dot{\alpha}(s)\,ds$$

The optimality condn is then $\xi(s) \equiv 0$ (this should give the PMP). Note that this calcn also tells us how to do a gradient-descent optimization if α_0 is not optimal.

We proceed as before: the 1st variation of the objective is

$$\int_t^T h_y \dot{y} + h_\alpha \dot{\alpha}\,ds + g_y \dot{y}(T)$$

(note: y + α are vector-valued, so $h_y \dot{y}$ really means the inner product $\langle \nabla_y h, \dot{y}\rangle$, etc).

Differentiating the state eqn gives

$$\dot{y}_s = f_y \dot{y} + f_\alpha \dot{\alpha}, \qquad \dot{y}(t) = 0.$$

Making an inspired choice, we let $p_0(s)$ solve

$$\frac{d}{ds} p_0 = -f_y(y_0(s), \alpha_0(s))\, p_0(s) - h_y(y_0(s), \alpha_0(s))$$

for $t<s<T$, with final-time condition

$$p_0(T) = g_y(y_0(T)).$$

(Motivation: at an optimal control we expect $p_0(s)$ to be the $\pi(s)$ of the PMP, and our eqn for p_0 says $\partial_s p_0 = -\nabla_y H$.)

Recall our goal: we want to write the 1st varn of the objective

$$\int_t^T h_y \dot{y} + h_\alpha \dot{\alpha}\, ds + g_y \dot{y}(T)$$

in the form $\int_t^T \xi(s) \dot{\alpha}(s)\, ds$. The term $h_\alpha \dot{\alpha}$ already has this form; the point of p_0 is to help us rewrite the others. In fact (writing p for p_0, y for y_0, α for α_0)

$$\int_t^T p_s \dot{y}\, ds = \int_t^T -f_y p \dot{y} - h_y \dot{y}\, ds$$

by the ode for p;

$$\text{LHS} = g_y(y(T))\dot{y}(T) - \underbrace{\cancel{p(t)\dot{y}(t)}}_{0} - \int_t^T p \dot{y}_s\, ds$$

$$\text{RHS} = \int_t^T -\dot{y}_s p + f_\alpha \dot{\alpha} p - h_y \dot{y}\, ds$$

Thus "LHS = RHS" becomes

$$g_y(y(T))\,\dot{y}(T) = \int_t^T f_\alpha \dot{\alpha} p - h_y \dot{y}\, ds.$$

We conclude that 1st varn of the objective is

$$\left(g_y \dot{y}(T) + \int_t^T h_y \dot{y}\, ds \right) + \int_t^T h_\alpha \dot{\alpha}\, ds.$$

$$= \int_t^T (f_\alpha \dot{\alpha} p + h_\alpha \dot{\alpha})\, ds$$

$$= \int_t^T \xi(s)\,\dot{\alpha}(s) \quad \text{with } \xi = f_\alpha p + h_\alpha.$$

To recover the PMP, we observe that if $\alpha_0(s)$ is optimal then $\xi(s) \equiv 0$, so

$$h_\alpha + f_\alpha p = 0$$

ie $\alpha(s)$ satisfies the 1st order optimality condition for achieving the optimum in

$$H(p, x) = \max_\alpha \left\{ f(x,\alpha)\cdot p + h(x,\alpha) \right\}$$

This relation, combined with the ODE for p and the state eqn, are precisely the PMP.

(As mentioned earlier, we have not *proved* existence of an optimal control; this requires some structural conditions. A deeper treatment of the PMP gives suff't as well as necessary cond's for optimality; see refs mentioned on pp 7.1 + 7.2 for more in this direction.)

Suggested exercises:

(1) Use the method of characteristics for the HJB eqn $u_t + H(x, \nabla u) = 0$ to show that when $y(s)$ + $\pi(s)$ solve the ODE's of the PMP, we have $\nabla u(y(s), s) = \pi(s)$ for all $t < s < T$.

(2) Complete the example begun on pp 7.7-7.8. The value function for that optimal control problem can be determined by solving the HJ eqn., using Exercise 4 of Lecture 6 (ie this example is a special case of the "linear quadratic regulator"). Is this fact of any use in solving the PMP?

(3) Our other favorite class of examples was the "minimum arrival time" problem

$$\min_{\alpha(s)\in A} \{ 1^{st} \text{ time } y(s) \text{ reaches } \Gamma \}$$

where Γ is a specified "target set" and the state eqn is

$$dy/ds = f(y(s), \alpha(s)), \quad y(0) = x .$$

What should the Pontryagin Max Prin say in this case? (Check the case $x \in D$, $\Gamma = \partial D$, $f(y(s), \alpha(s)) = \alpha(s)$, $A = \{|\alpha| \le 1\}$ to be sure your answer is reasonable.)

(4) Consider the "optimal consumption" example we discussed in Lecture 6, with $\rho = 0$ for simplicity:

$$\max_{a(s)} \int_t^T a^{\gamma}(s)\, ds .$$

with state eqn

$$\frac{dy}{ds} = ry - a \qquad t < s < T, \quad y(t) = x,$$

(Here $0 < \gamma < 1$ and $r > 0$ are fixed constants.)

(i) What ode's and end conditions does the PMP give in this case? How is the optimal control related to $\pi(s)$?

(ii) Show that $a(s) = Ce^{rs/(1-\gamma)}$ for some constant C.

[To find a complete, explicit soln this way requires solving a nonlinear eqn for the value of $\pi(t)$ that gives $\pi(T)=0$. This is messy, so I'm not suggesting you do it by hand.].

(5) In part B of these notes, our example of sensitivity analysis used a pde with a Dirichlet boundary condition. Consider a similar example with a Neumann boundary condition: given $a(x)$, let u_a solve

$$-\Delta u + a(x) u = 1 \quad \text{in } D$$
$$\partial u/\partial n = 0 \quad \text{at } \partial D$$

(There is a unique soln if, for example, $a(x) > 0$.). Given $a_0(x) > 0$ and the assoc u_0, find the "gradient" of $F[a] = \int_D |u_a - g|^2$, ie find $\xi(x)$ st

$$\frac{d}{dt}\Big|_{t=0} F[a_0 + t\dot{a}] = \int_D \dot{a}(x)\,\xi(x)\,dx$$

(6) Another sensitivity analysis question: suppose this time u_a is defined by solving

$$-\mathrm{div}\,(a(x)\nabla u) = 1 \quad \text{in } D$$
$$u = 0 \quad \text{at } \partial D$$

(I suppose here that $a(x) > 0$). Given $a_0(x) > 0$ and the assoc u_0, find the "gradient" of $F[a] = \int_D |u_a - g|^2 dx$.