Appendix A – Hanson 2007 Book

User 1006 • 17 Sep 2024

Very entry-level book (e.g., no measurability) focusing on the applied aspects.

Appendix A

This appendix collects results from deterministic control theory. Appendix B sketches background material in probability.

Deterministic Control via PMP

Problem formulation (Bolza form):

$V[\boldsymbol{X}, \boldsymbol{U}, t_f](\boldsymbol{x}_0, t_0) \equiv \int_{t_0}^{t_f} C(\boldsymbol{X}(t), \boldsymbol{U}(t), t)dt + S(\boldsymbol{X}(t_f), t_f)$

$\frac{d\boldsymbol{X}}{dt}(t) = \boldsymbol{f}(\boldsymbol{X}(t), \boldsymbol{U}(t), t);\ \ \boldsymbol{X}(t_0) = \boldsymbol{x}_0$

Here $\boldsymbol{X}(t)$ is a n_x-dim state vector; $\boldsymbol{U}(t)$ is a n_u-dim control vector.

The book calls the control space $\mathcal{U}$ to be the whole space $\mathbb{R}^{n_u}$ if the problem is unconstrained. I think it's a bit loose, normally we be explicit and say $\mathcal{U}$ is the set of processes valued in U, where U is this Euclidean space.

cost function
$v^*(\boldsymbol{x}_0, t_0) = \max_{\boldsymbol{U} \in \mathcal{U}} [V[\boldsymbol{X}, \boldsymbol{U}, t_f](\boldsymbol{x}_0, t_0)]$

We would search over all critical points: **regular points** (gradient = 0), **boundary points** of the control domain, and **singular points** or other irregular points. Here singular points probably is referring to that in ode: consider $\frac{y'}{x-1} = 0$. Then $x = 1$ is a singular point.

Hamiltonian
$\mathcal{H}(\boldsymbol{X}(t), \boldsymbol{U}(t), \boldsymbol{\lambda}(t), t) \equiv C(\boldsymbol{X}(t), \boldsymbol{U}(t), t) + \boldsymbol{\lambda}^T(t)\boldsymbol{f}(\boldsymbol{X}(t), \boldsymbol{U}(t), t)$

where $\boldsymbol{\lambda}(t)$ is of n_x-dim, called the Lagrange multiplier/adjoint state/co-state. It provides the objective extension for including the state dynamics.

PMP, regular control, interior point optimum
If Hamiltonian $\mathcal{H}$ is C^1 in $(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\lambda})$, then necessary conditions for an interior point optimum, in terms of the Hamiltonian $\mathcal{H}$ along the optimal trajectories, are:

1. $\frac{d\boldsymbol{X}^*}{dt}(t) = (\frac{\partial \mathcal{H}}{\partial \boldsymbol{\lambda}})^* \equiv (\frac{\partial \mathcal{H}}{\partial \boldsymbol{\lambda}})(\boldsymbol{X}^*(t), \boldsymbol{U}^*(t), \boldsymbol{\lambda}^*(t), t) = \boldsymbol{f}(\boldsymbol{X}^*(t), \boldsymbol{U}^*(t), t)$
2. $-\frac{d\boldsymbol{\lambda}^*}{dt}(t) = (\frac{\partial \mathcal{H}}{\partial \boldsymbol{x}})^*$
3. $\boldsymbol{0} = (\frac{\partial \mathcal{H}}{\partial \boldsymbol{u}})^*$

Note that these do not necessarily apply to boundary or singular points of the control. The associated final conditions were listed in the book, depending on whether terminal time/state are fixed or flexible.

Long proof (including boundary conditions) via formal use of calculus of variations.

The necessary conditions become sufficient if:

1. $\mathcal{H}$ is C^2 in $\boldsymbol{u}$
2. the Hessian matrix $\mathcal{H}^*_{\boldsymbol{uu}} \equiv \nabla_{\boldsymbol{u}}[\nabla^T_{\boldsymbol{u}}[\mathcal{H}]](\boldsymbol{X}^*(t), \boldsymbol{U}^*(t), \boldsymbol{\lambda}^*(t), t)$ is neg.def (for maximum) or pos.def (for minimum). Or, $\mathcal{H}$ is concave (for maximum) or convex (for minimum) at the regular point.

There are cases when this principle cannot be used. For example, if $\mathcal{H}$ is linear in control, then it is a singular control problem and need to use more basi principles.

Book gave an example of regular control problem with explicit solutions: determining optimal consumption rate. Also, illustrated why second-order condition is needed – if agent is risk-seeking, then there is no interior optimum, agents would like to use bang control and consume at the maximum rate since they do not like consumption smoothing.

(more) Basic Optimum Principle
common situations: optima are at some interior points where the derivatives in Hamilton's equations do not exist, or, optima at boundaries of control set, or linear control problem.

Optimum Principles Necessary conditions for maximum principle is that $\mathcal{H}$ is maximized: $\mathcal{H}^* = \mathcal{H}(\boldsymbol{X}^*(t), \boldsymbol{U}^*(t), \boldsymbol{\lambda}^*(t), t) \geq \mathcal{H}(\boldsymbol{X}^*(t), \boldsymbol{U}(t), \boldsymbol{\lambda}^*(t), t)$.

If looking for minimum, we minimize Hamiltonian.

The book offered another expression of the maximum principle, as $\Delta_u \mathcal{H}(\boldsymbol{X}^*(t), \boldsymbol{U}^*(t), \boldsymbol{\lambda}^*(t), t) \leq 0$, where we define:

$\Delta_u \mathcal{H}(\boldsymbol{X}^*(t), \boldsymbol{U}^*(t), \boldsymbol{\lambda}^*(t), t) \equiv \mathcal{H}(\boldsymbol{X}^*(t), \boldsymbol{U}^*(t) + \delta\boldsymbol{U}(t), \boldsymbol{\lambda}^*(t), t) - \mathcal{H}(\boldsymbol{X}^*(t), \boldsymbol{U}^*(t), \boldsymbol{\lambda}^*(t), t)$

But I find it not very necessary, as long as we remember the principle is to maximize the Hamiltonian, whether it is at interior point or not.

Book then gave one example of linear control, and showed the optimal control is bang-bang; and an example of singular control (a bit convoluted), where the control is bang-singular-bang. Although I find the singular control example illuminating, the explanation at places could be clearer imo, and I find the analysis method hard to generalize to more complicated problems.

Linear Quadratic Models

There is a motivating example for the scalar case presented first, then moving to the multi-dim case. I'll summarize this case only since the conditions will be derived again using DP later.

The problem is:

$V[\boldsymbol{X}, \boldsymbol{U}, t_f](\boldsymbol{x_0}, t_0) = \frac{1}{2}\int_{t_0}^{t_f}[\boldsymbol{X}^T(t)Q(t)\boldsymbol{X}(t) + \boldsymbol{U}^T(t)R(t)\boldsymbol{U}(t)]dt + \frac{1}{2}\boldsymbol{X}^T(t_f)S_f(t_f)\boldsymbol{X}(t_f)$

where the state LOM is:

$\dot{\boldsymbol{X}}(t) = A(t)\boldsymbol{X}(t) + B(t)\boldsymbol{U}(t);\ \boldsymbol{X}(t_0) = \boldsymbol{x_0}$

Setting up the Hamiltonian:
$\mathcal{H}(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\lambda}, t) = \frac{1}{2}(\boldsymbol{x}^T Q(t)\boldsymbol{x} + \boldsymbol{u}^T R(t)\boldsymbol{u}) + \boldsymbol{\lambda}^T(A(t)\boldsymbol{x} + B(t)\boldsymbol{u})$, the optimality conditions are:

1. $\dot{\boldsymbol{X}}^*(t) = \mathcal{H}^*_{\boldsymbol{\lambda}} = A(t)\boldsymbol{X}^*(t) + B(t)\boldsymbol{U}^*(t)$
2. $\dot{\boldsymbol{\lambda}}^*(t) = -\mathcal{H}^*_{\boldsymbol{x}} = -Q(t)\boldsymbol{X}^*(t) - A^T(t)\boldsymbol{\lambda}^*(t)$
3. $\boldsymbol{0} = \mathcal{H}^*_{\boldsymbol{u}} = R(t)\boldsymbol{U}^*(t) + B^T(t)\boldsymbol{\lambda}^*(t)$

From 3, we get $\boldsymbol{U}^*(t) = -R^{-1}(t)B^T(t)\boldsymbol{\lambda}^*(t)$. We will now solve for $\boldsymbol{\lambda}^*(t)$.

Denote $\boldsymbol{Z}(t) \equiv \begin{bmatrix} \boldsymbol{X}^*(t) \\ \boldsymbol{\lambda}^*(t) \end{bmatrix}$, collect the conditions as:

$\dot{\boldsymbol{Z}}(t) = M(t)\boldsymbol{Z}(t)$, where:

$M(t) \equiv \begin{bmatrix} A(t) & -B(t)R^{-1}(t)B^T(t) \\ -Q(t) & -A^T \end{bmatrix}$.

Given the linear structure, assume the solution takes the form $\boldsymbol{Z}(t) = Ce^{\mu t}\boldsymbol{\zeta}$, where this $\boldsymbol{\zeta}$ is a vector of dimension $(2n_x)$.

Plugging in, we find the equations become $M(t)\boldsymbol{\zeta} = \mu\boldsymbol{\zeta}$. Therefore, the μ and the ζ are eigen-values and eigen-vectors.

The general solutions is of the form: $\boldsymbol{Z}(t) = \sum_{k=1}^{2n_x} C_k e^{\mu_k t}\zeta_k$, and we shall use the initial and terminal/boundary conditions to pin down the constants $\boldsymbol{C}$.

Then we really have solved for everything, and we can derive other terms of interest, like the explicit form of the linear feedback rule.

There is an alternative approach, which is to start by assuming $\boldsymbol{\lambda}^*(t) = J(t)\boldsymbol{X}^*(t)$, and then differentiate this equation wrt time and match coefficients to get a matrix Riccati equation for $J(t)$.

Note that this $J(t)$ will reappear in DP approach, as the value function turns out to have the form $v^*(\boldsymbol{x}, t) = \frac{1}{2}\boldsymbol{x}^T J(t)\boldsymbol{x}$. This is to be expected as λ in PMP is the Jacobian of the value function wrt the state variable x.

Deterministic Dynamic Programming

now, initial state is arbitrary: $\boldsymbol{X}(t) = \boldsymbol{x}$.

Bellman's principle:
$v^*(\boldsymbol{x}, t) = \min_{\boldsymbol{U}(t,t+\Delta t]} [\int_t^{t+\Delta t} C(\boldsymbol{X}(s), \boldsymbol{U}(s), s)ds + v^*(\boldsymbol{x} + \Delta\boldsymbol{X}(t), t + \Delta t)]$

infinitesimal version of Bellman – HJB:
If $v^*(\boldsymbol{x}, t)$ is differentiable in $\boldsymbol{x}$ and t, then first define pseudo-Hamiltonian:

$\mathcal{H}(\boldsymbol{x}, \boldsymbol{u}, t) \equiv C(\boldsymbol{x}, \boldsymbol{u}, t) + \nabla_{\boldsymbol{x}}^T[v^*](\boldsymbol{x}, t) \cdot \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}, t)$

Then HJB is written as:
$0 = v_t^*(\boldsymbol{x}, t) + \min_{\boldsymbol{u}} [\mathcal{H}(\boldsymbol{x}, \boldsymbol{u}, t)] \equiv v_t^*(\boldsymbol{x}, t) + \mathcal{H}^*(\boldsymbol{x}, t)$

Book solved the same LQ problem using DP approach. Guess the functional form of the value function, and derived the matrix Riccati equation that is consistent with that derived via PMP.

recall the problem was maximizing:
$V[\boldsymbol{X}, \boldsymbol{U}, t_f](\boldsymbol{x}, t) = \frac{1}{2}\int_t^{t_f}[\boldsymbol{X}^T(s)Q(s)\boldsymbol{X}(s) + \boldsymbol{U}^T(s)R(s)\boldsymbol{U}(s)]ds + \frac{1}{2}\boldsymbol{X}^T(t_f)S_f(t_f)\boldsymbol{X}(t_f)$

subject to $\dot{\boldsymbol{X}}(s) = A(s)\boldsymbol{X}(s) + B(s)\boldsymbol{U}(s);\ \boldsymbol{X}(t) = \boldsymbol{x}$

Thus the Hamiltonian is:
$\mathcal{H}(\boldsymbol{x}, \boldsymbol{u}, t) = \frac{1}{2}(\boldsymbol{x}^T Q(t)\boldsymbol{x} + \boldsymbol{u}^T R(t)\boldsymbol{u}) + \nabla_{\boldsymbol{x}}^T[v^*](\boldsymbol{x}, t)(A(t)\boldsymbol{x} + B(t)\boldsymbol{u})$

Now we guess $v_t^*(\boldsymbol{x}, t) = \frac{1}{2}\boldsymbol{x}^T J(t)\boldsymbol{x}$, where $J(t)$ is symmetric matrix.

Then, $\nabla_{\boldsymbol{x}}[v^*](\boldsymbol{x},t) = J(t)\boldsymbol{x}$, and upon plugging into HJB we get:

$\boldsymbol{u}^*(\boldsymbol{x},t) = -R^{-1}(t)B^T(t)J(t)\boldsymbol{x}$

Note that this is the linear feedback control derived earlier.

Now, $v_t^*(\boldsymbol{x},t) = \boldsymbol{x}^T J'\boldsymbol{x}$, and plugging the optimal control into Hamiltonian, we see from the HJB that we will have a matrix Riccati equation for $J(t)$, with a boundary condition, and they are identical to the ones derived via PMP.

control of PDE driven dynamics

This is about distributed parameter systems. Better to learn from a more systematic source on this topic.

This concludes the review of deterministic control part.