Pontryagins Maximum Principle in a Nutshell

Moritz Diehl

Overview

- Derivation 1: Hamilton-Jacobi-Bellman equation
- Derivation 2: Calculus of Variations
- Properties of Euler-Lagrange Equations
- Boundary Value Problem (BVP) Formulation
- Numerical Solution of BVP
- Discrete Time Pontryagin Principle

Continuous Time Optimal Control

Regard simplified optimal control problem:

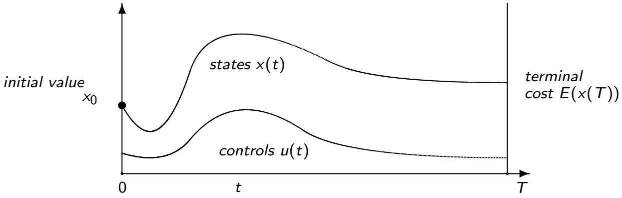

$$\underset{x(\cdot),u(\cdot)}{\text{minimize}} \quad \int_0^T L(x(t), u(t))\, dt + E\,(x(T))$$

subject to

$$\begin{aligned} x(0) - x_0 &= 0, & & \text{(fixed initial value)} \\ \dot{x}(t) - f(x(t), u(t)) &= 0, \quad t \in [0, T]. & & \text{(ODE model)} \end{aligned}$$

Hamilton-Jacobi-Bellman equation

Recall:

- Hamilton-Jacobi-Bellman Equation:
  $-\frac{\partial J}{\partial t}(x,t) = \min_u H(x, \nabla J(x,t), u)$
- with Hamiltonian function
  $H(x, \lambda, u) := L(x,u) + \lambda^T f(x,u)$
- and terminal condition $J(x, T) = E(x)$.

Pontryagin's Maximum Principle

OBSERVATION: In HJB, optimal controls

$$u^*(x,t) = \arg\min_u H(x, \nabla_x J(x,t), u)$$

depend only on derivative $\nabla_x J(x,t)$, not on J itself!

IDEA: Introduce **adjoint variables**
$\lambda(t) \quad \hat{=} \quad \frac{\partial J}{\partial x}(x(t),t)^T \in \mathbb{R}^{n_x}$ and get controls from
Pontryagin's Maximum Principle (historical name)

$$u^*(x,\lambda) = \arg\min_u H(x, \lambda, u)$$

QUESTION: How to obtain $\lambda(t)$?

(Differentiation Lemma)

We want to differentiate optimal solution that depends on parameters $y = (x, \lambda)$. How can we do that easiest?

LEMMA: If $H^*(y) = \min_u H(y, u)$

then $\frac{\partial H^*}{\partial y}(y) = \frac{\partial H}{\partial y}(y, u^*)$

with $u^* = \arg\min_u H(y, u)$

PROOF: $\frac{\partial H^*}{\partial y}(y) = \frac{\partial H}{\partial y}(y, u^*) + \underbrace{\frac{\partial H}{\partial u}(y, u^*)}_{=0} \frac{\partial u^*}{\partial y}(y)$

due to the first order optimality condition. $\square$

(Lemma can be extended to constrained problems, using partial derivatives of Lagrangian.)

Adjoint Differential Equation

- Differentiate HJB Equation

$$-\frac{\partial J}{\partial t}(x,t) = \min_u H(x, \nabla J(x,t), u) = H(x, \nabla_x J(x,t), u^*) =$$

with respect to x and obtain:

$$-\frac{\partial \nabla J^T}{\partial t} = \frac{\partial H}{\partial x} + \underbrace{\frac{\partial H}{\partial \lambda}}_{=f(x,u^*)^T} \quad \nabla_x^2 J(x,t), u^*)$$

or equivalently

$$\frac{\partial \nabla J}{\partial t} + \nabla_x^2 J(x,t), u^*) f(x,u^*) = \underbrace{\frac{\partial}{\partial t} \nabla_x J(x,t)}_{=\dot{\lambda}(t)} = -\nabla_x H(x, \lambda, u^*)$$

Terminal Condition

- Likewise, differentiate $J(x, T) = E(x)$ and obtain terminal condition

$$\lambda(T) = \nabla E(x(T)).$$

Necessary Optimality Conditions

Summarize optimality conditions as **boundary value problem**:

$$
\begin{array}{rcllr}
x(0) & = & x_0, & & \text{(initial value)} \\
\dot{x}(t) & = & f(x(t), u^*(t)) & t \in [0, T], & \text{(ODE model)} \\
\dot{\lambda}(t) & = & -\frac{\partial H}{\partial x}(x(t), u^*(t), \lambda(t))^T, & t \in [0, T], & \text{(adjoint equations)} \\
u^*(t) & = & \arg\min\limits_{u} H(x(t), u, \lambda(t)), & t \in [0, T], & \text{(minimum principle)} \\
\lambda(T) & = & \nabla E(x(T)). & & \text{(adjoint final value).}
\end{array}
$$

Overview

- Derivation 1: Hamilton-Jacobi-Bellman equation
- **Derivation 2: Calculus of Variations**
- Properties of Euler-Lagrange Equations
- Boundary Value Problem (BVP) Formulation
- Numerical Solution of BVP

Alternative Derivation

Regard infinite optimization problem:

$$\underset{x(\cdot),u(\cdot)}{\text{minimize}} \quad \int_0^T L(x(t),u(t))\,dt + E\left(x(T)\right)$$

subject to

$$\begin{aligned} x(0) - x_0 &= 0, && && \text{(fixed initial value)} \\ \dot{x}(t) - f(x(t),u(t)) &= 0, && t \in [0,T]. && \text{(ODE model)} \end{aligned}$$

Introduce Lagrangian multipliers λ and "Lagrangian functional"

$$\mathcal{L}(x(\cdot),u(\cdot),\lambda(\cdot)) = \int_0^T L(x,u) - \lambda^T(\dot{x} - f(x,u))dt + E(x(T))$$

.

Infinitesimal Variations

Abbreviate using the Hamiltonian $H(x, \lambda, u)$

$$\mathcal{L}(x(\cdot), u(\cdot), \lambda(\cdot)) = \int_0^T H(x, \lambda, u) - \lambda^T \dot{x} dt + E(x(T))$$

Regard infinitesimal variation of $\mathcal{L}(x(\cdot), u(\cdot), \lambda(\cdot))$ with perturbation $\delta x(t)$

$$\delta \mathcal{L} = \int_0^T \frac{\partial H}{\partial x} \delta x - \lambda^T \delta \dot{x} dt + \frac{\partial E}{\partial x} \delta x(T)$$

which by partial integration yields:

$$\delta \mathcal{L} = \int_0^T (\nabla_x H + \dot{\lambda})^T \delta x - \frac{d}{dt}(\lambda^T \delta x) dt + \frac{\partial E}{\partial x} \delta x(T)$$

Infinitesimal Variations (contd.)

Using the fact that $\delta x(0) = 0$ and requiring $\delta \mathcal{L} = 0$ yields

$$0 = \int_0^T (\nabla_x H + \dot{\lambda})^T \delta x dt + (\nabla_x E - \lambda(T))^T \delta x(T)$$

which implies, for arbitrary variations,

$$\dot{\lambda} = -\nabla_x H(x(t), \lambda(t), u^*(t))$$

and

$$\lambda(T) = \nabla_x E(x(T))$$

Thus, calculus of variations leads to the same adjoint differential equations as differentiation of HJB!

Overview

- Derivation 1: Hamilton-Jacobi-Bellman equation
- Derivation 2: Calculus of Variations
- **Properties of Euler-Lagrange Equations**
- Boundary Value Problem (BVP) Formulation
- Numerical Solution of BVP

How to obtain explicit expression for controls?

- In simplest case,

$$u^*(t) = \arg\min_u H(x(t), \lambda(t), u)$$

  is defined by

$$\frac{\partial H}{\partial u}(x(t), \lambda(t), u^*(t)) = 0$$

- In presence of path constraints, expression for $u^*(t)$ changes whenever active constraints change. This leads to state dependent switches.
- If minimum of Hamiltonian locally not unique, "singular arcs" occur. Treatment needs higher order derivatives of H.

Nice Case: Example

Regard $L(x,u) = \frac{1}{2}(x^TQx + u^TRu)$ with invertible R and $f(x,u) = Ax + Bu$. Then

$$H(x,\lambda,u) = \frac{1}{2}(x^TQx + u^TRu) + \lambda^T(Ax + Bu).$$

and

$$\frac{\partial H}{\partial u} = u^TR + \lambda^TB.$$

Thus, $\frac{\partial H}{\partial u} = 0$ implies that

$$u^* = -R^{-1}B^T\lambda$$

Singular Arcs

But what if the relation

$$\frac{\partial H}{\partial u}(x(t), \lambda(t), u^*) = 0$$

is not invertible w.r.t. to u^*?
This e.g. occurs if $L(x, u)$ is independent of u and $f(x, u)$ is linear in u.
Singular arcs are due to the fact that only the integral of controls influences the states, and "singular" perturbations (that go up and down quickly) do not matter in the objective.

Remedy for Singular Arcs

"What is zero should also have zero derivative".
Therefore, we differentiate totally w.r.t. to time

$$\frac{d}{dt}\frac{\partial H}{\partial u}(x(t), \lambda(t), u^*) = 0$$

i.e.

$$\frac{\partial}{\partial x}\frac{\partial H}{\partial u}\underbrace{\dot{x}}_{=f(x,u)} + \frac{\partial}{\partial \lambda}\frac{\partial H}{\partial u}\underbrace{\dot{\lambda}}_{=-\frac{\partial H}{\partial x}} = 0$$

If this still does not allow to find u^* explicitly, differentiate even further ...

Singular Arc: Example

Regard $L(x, u) = x^T Qx$ and $f(x, u) = Ax + Bu$. Then

$$H(x, \lambda, u) = \frac{1}{2}x^T Qx + \lambda^T (Ax + Bu).$$

and

$$\frac{\partial H}{\partial u} = \lambda^T B$$

This is not invertible w.r.t. to u^*!

Singular Arc: Example (contd.)

Once more differentiating yields:

$$\frac{d}{dt}\frac{\partial H}{\partial u} = \dot{\lambda}^T B = -\frac{\partial H}{\partial x} B = -(x^T Q + \lambda^T A)B$$

Once more differentiating yields:

$$\frac{d}{dt}\frac{d}{dt}\frac{\partial H}{\partial u} = -\dot{x}^T QB - \dot{\lambda}^T AB = -(Ax+Bu)^T QB + (x^T Q + \lambda^T A)AB$$

Setting this to zero finally yields the feedback law

$$u^* = (B^T QB)^{-1} B^T \left((A^T Q - QA)x + A^T A^T \lambda \right)$$

This is only applicable on singular arcs.

Euler Lagrange Differential Equations

Note that

$$\frac{\partial}{\partial \lambda} H(x, \lambda, u) = f(x, u)$$

Thus,

$$\frac{d}{dt}\begin{bmatrix} x \\ \lambda \end{bmatrix} = \begin{bmatrix} \frac{\partial H}{\partial \lambda} \\ -\frac{\partial H}{\partial x} \end{bmatrix}$$

is a Hamiltonian system. Volume in (x, λ) is preserved. But this also means that if the dynamics of x is very stable i.e. contracting than the dynamics of λ must be expanding.

Overview

- Derivation 1: Hamilton-Jacobi-Bellman equation
- Derivation 2: Calculus of Variations
- Properties of Euler-Lagrange Equations
- **Boundary Value Problem (BVP) Formulation**
- Numerical Solution of BVP

Boundary Value Problem (BVP)

Can summarize the BVP

$$
\begin{array}{rcllr}
x(0) & = & x_0, & & \text{(initial value)} \\
\dot{x}(t) & = & f(x(t), u^*(t)) & t \in [0, T], & \text{(ODE model)} \\
-\dot{\lambda}(t) & = & \frac{\partial H}{\partial x}(x(t), \lambda(t), u^*(t))^T, & t \in [0, T], & \text{(adjoint equations)} \\
u^*(t) & = & \arg\min_u H(x(t), \lambda(t), u), & t \in [0, T], & \text{(minimum principle)} \\
\lambda(T) & = & \frac{\partial E}{\partial x}(x(T))^T. & & \text{(adjoint final value).}
\end{array}
$$

by using $y = (x, \lambda)$ and substituting u^* explicitly as

$$
\begin{array}{rcllr}
0 & = & r(y(0), y(T)), & & \text{(boundary conditions)} \\
\dot{y}(t) & = & \tilde{f}(y(t)) & t \in [0, T], & \text{(ODE model)}
\end{array}
$$

BVP analysis

The BVP

$$
\begin{aligned}
0 &= r(y(0), y(T)), & & \text{(boundary conditions)} \\
\dot{y}(t) &= \tilde{f}(y(t)) & t \in [0, T], & \quad \text{(ODE model)}
\end{aligned}
$$

has $2n_x$ differential equations $\dot{y} = \tilde{f}$, and $2n_x$ boundary conditions r. It is therefore (usually) well-defined.
But how to solve a BVP?

Overview

- Derivation 1: Hamilton-Jacobi-Bellman equation
- Derivation 2: Calculus of Variations
- Properties of Euler-Lagrange Equations
- Boundary Value Problem (BVP) Formulation
- **Numerical Solution of BVP**
  - single shooting
  - collocation

Single Shooting

Guess initial value for y_0. Use numerical integration to obtain trajectory as function $y(t; y_0)$ of y_0.

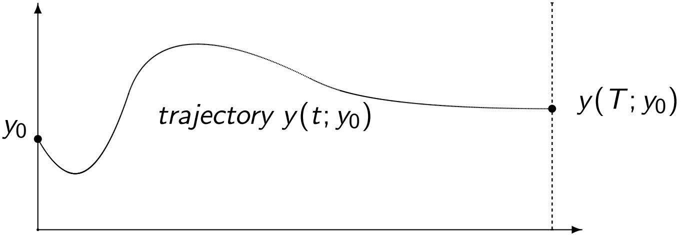

Obtain in particular terminal value $y(T; y_0)$.

Single Shooting (contd.)

The only remaining equation is

$$\underbrace{r(y_0, y(T; y_0))}_{=F(y_0)} = 0$$

which might or might not be satisfied for the guess y_0.
Fortunately, r has as many components as y_0, so we can apply Newton's method for root finding of

$$F(y_0) = 0$$

which iterates

$$y_0^{k+1} = y_0^k - \frac{\partial F}{\partial y_0}(y_0^k) F(y_0^k)$$

Attention: to evaluate $\frac{\partial F}{\partial y_0}(y_0^k)$ we have to compute ODE sensitivities.

Collocation (Sketch)

- Discretize states on grid with node values $s_i \approx y(t_i)$.
- Replace infinite ODE

$$0 = \dot{y}(t) - \tilde{f}(y(t)), \quad t \in [0, T]$$

by finitely many equality constraints

$$\begin{array}{rrcl} & c_i(s_i, s_{i+1}) & = & 0, \quad i = 0, \dots, N-1, \\ \text{e.g.} & c_i(s_i, s_{i+1}) & := & \frac{s_{i+1} - s_i}{t_{i+1} - t_i} - \tilde{f}\left(\frac{s_i + s_{i+1}}{2}\right) \end{array}$$

Nonlinear Equation in Collocation

After discretization, obtain large scale, but sparse nonlinear equation system:

$$
\begin{aligned}
r(s_0, s_N) &= 0, & & \text{(boundary conditions)} \\
c_i(s_i, s_{i+1}) &= 0, \quad i = 0, \ldots, N-1, & & \text{(discretized ODE model)}
\end{aligned}
$$

Solve again with Newton's method. Exploit sparsity in linear system setup and solution.

Discrete Time Optimal Control Problem

$$\underset{s,q}{\text{minimize}} \quad \sum_{i=0}^{N-1} l_i(s_i, q_i) \;+\; E\,(s_N)$$

subject to

$$\begin{aligned}
s_0 - x_0 &= 0, & & \text{(initial value)}\\
s_{i+1} - f_i(s_i, q_i) &= 0, \quad i = 0, \ldots, N-1, & & \text{(discrete system)}\\
h_i(s_i, q_i) &\geq 0, \quad i = 0, \ldots, N, & & \text{(path constraints)}\\
r\,(s_N) &\geq 0. & & \text{(terminal constraints)}
\end{aligned}$$

Can arise also from direct multiple shooting parameterization of continous optimal control problem. This NLP can be solved by SQP or Constrained Gauss-Newton method.

Optimality of Discretized Optimal Control

For simplicity, drop all inequalities, regard only:

$$\underset{s,q}{\text{minimize}} \quad \sum_{i=0}^{N-1} l_i(s_i, q_i) \; + \; E(s_N)$$

subject to

$$\begin{aligned} s_0 - x_0 &= 0, & & \text{(initial value)} \\ s_{i+1} - f_i(s_i, q_i) &= 0, \quad i = 0, \ldots, N-1, & & \text{(discrete system)} \end{aligned}$$

What are KKT optimality conditions of this discretized optimal control problem?

Discrete Optimality Conditions 1

Procedure:

- Introduce multiplier vectors $\lambda_0, \ldots, \lambda_N$ for all dynamic state constraints.
- Formulate Lagrangian

$$\mathcal{L}(s, q, \lambda) = E(s_N) + (s_0 - x_0)^T \lambda_0 + \sum_{i=0}^{N-1} l_i(s_i, q_i) - (s_{i+1} - f_i(s_i, q_i))^T \lambda_{i+1}$$

- Compute

$$\nabla_{s_i} \mathcal{L} \quad \text{and} \quad \nabla_{q_i} \mathcal{L},$$

which must be zero for optimal solution.

Discrete Optimality Conditions 2

Obtain

1. $\nabla_{s_i}\mathcal{L} = -\lambda_i + \nabla_{s_i} l_i(s_i, q_i) + \nabla_{s_i} f_i(s_i, q_i)\lambda_{i+1} = 0$
   $(i = 0, \ldots, N-1)$
2. $\nabla_{s_N}\mathcal{L} = -\lambda_N + \nabla_{s_N} E(s_N)$
3. $\nabla_{q_i}\mathcal{L} = \nabla_{q_i} l_i(s_i, q_i) + \nabla_{q_i} f_i(s_i, q_i)\lambda_{i+1} = 0 \quad (i = 0, \ldots, N-1)$
4. $\nabla_{\lambda_0}\mathcal{L} = s_0 - x_0$
5. $\nabla_{\lambda_{i+1}}\mathcal{L} = s_{i+1} - f_i(s_i, q_i) = 0 \quad (i = 0, \ldots, N-1)$

These conditions can be simplified by introducing the discrete time Hamiltonian:

$$H_i(s_i, q_i, \lambda_{i+1}) = l_i(s_i, q_i) + f_i(s_i, q_i)^T \lambda_{i+1}$$

as follows...

Discrete Pontryagin Principle

The KKT conditions are now equivalent to:

1. $\lambda_i = \nabla_{s_i} H_i(s_i, q_i, \lambda_{i+1}) \quad (i = 0, \ldots, N-1)$ (adjoint equation)
2. $\lambda_N = \nabla_{s_N} E(s_N)$ (terminal condition on adjoints)
3. $\nabla_{q_i} H_i(s_i, q_i, \lambda_{i+1}) = 0 \quad (i = 0, \ldots, N-1)$ (minimum principle)
4. $s_0 = x_0$ (initial condition)
5. $s_{i+1} = f_i(s_i, q_i) \quad (i = 0, \ldots, N-1)$ (system dynamics)

Summary

- Pontryagin's Maximum Principle can be derived by in two ways (HJB/Calc. of Variations)
- Controls must be explicitly derived. On singular arcs, need higher order derivatives.
- Boundary Value Problem is well posed but double ODE is usually unstable. Not easy to simulate forward.
- Solve BVP with single shooting, collocation, or multiple shooting.

Literature

- A. E. Bryson and Y. C. Ho: Applied Optimal Control, Hemisphere/Wiley, 1975.