

Lecture 19, Nov 26, 2025

Linear Quadratic Optimal Control

- Consider a spring-mass system with an input force, $u - kx_1 = m\ddot{x}_1$ and let $x_2 = \dot{x}_1$, so the system is described by $\dot{\mathbf{x}} = \begin{bmatrix} 0 & 1 \\ -k/m & 0 \end{bmatrix} \mathbf{x} + \begin{bmatrix} 0 \\ 1/m \end{bmatrix} u$
 - The energy of the system is $E = \frac{1}{2}kx_1^2 + \frac{1}{2}mx_2^2$
 - We can try to minimize the energy over time, $J = \int_0^\infty \frac{1}{2}kx_1^2(\tau) + \frac{1}{2}mx_2^2(\tau) + \frac{1}{2}u^2(\tau) d\tau$
 - Note we need to include the control input in the integral to make the problem well-defined, otherwise this will just cause infinite control inputs
- More generally, we can have any cost function quadratic in the system state and control input

Definition

For a stabilizable system $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$, let \mathcal{U} be the set of admissible control laws of the form $\mathbf{u}(t) = \phi(t)$, $t \geq 0$ such that $\phi : [0, \infty) \mapsto \mathbb{R}^m$ is continuous and asymptotically stabilizes the closed-loop system.

The *Linear Quadratic Optimal Control* problem is to find some $\phi \in \mathcal{U}$ which minimizes a quadratic cost function,

$$J(\mathbf{x}, \phi) = \int_0^\infty \mathbf{x}(t)^T \mathbf{Q} \mathbf{x}(t) + \phi(t)^T \mathbf{R} \phi(t) dt$$

where \mathbf{Q} is symmetric positive semidefinite, and \mathbf{R} is symmetric positive definite.

- Intuitively, the reason we choose this cost function is that it usually relates to the “energy” of the system, and penalizes the state being far from the origin and the control effort being large
- Let $L(\mathbf{x}, \mathbf{u}) = \mathbf{x}^T \mathbf{Q} \mathbf{x} + \mathbf{u}^T \mathbf{R} \mathbf{u}$ be the instantaneous cost for the current values of \mathbf{x}, \mathbf{u}
- Define the optimal value function $V(\mathbf{x}) = \inf_{\phi \in \mathcal{U}} J(\mathbf{x}, \phi) = \inf_{\phi \in \mathcal{U}} \int_0^\infty \mathbf{x}(t)^T \mathbf{Q} \mathbf{x}(t) + \mathbf{u}(t)^T \mathbf{R} \mathbf{u}(t) dt$ such that $\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t)$, $\mathbf{x}(0) = \mathbf{x}$, i.e. $V(\mathbf{x})$ is the lowest possible cost if we start with an initial state of \mathbf{x}
 - Note the infimum \inf denotes the greatest lower bound (we can't simply use \min since we don't know if one exists yet), e.g. $\inf_{n \in \mathbb{N}} \frac{1}{n} = 0$, but 0 is not in the set
- We can show that if an optimal control $\phi^* \in \mathcal{U}$ exists, i.e. $V(\mathbf{x}) = \min_{\phi^* \in \mathcal{U}} J(\mathbf{x}, \phi) = J(\mathbf{x}, \phi^*)$, then (V, ϕ^*) satisfies the *Hamilton-Jacobi-Bellman* (HJB) equation, $\min_{\mathbf{u} \in \mathbb{R}^m} \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \cdot (\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}) + \mathbf{L}(\mathbf{x}, \mathbf{u}) = 0$
 - Intuitively, since the first term is essentially $\frac{\partial V}{\partial t}$, this means that the cost does not change instantaneously
 - Let $\tau > 0$ and suppose ϕ^* exists, then $V(\mathbf{x}) = \int_0^\tau L(\mathbf{x}(t), \phi^*(t)) dt + \int_\tau^\infty L(\mathbf{x}(t), \phi^*(t)) dt$
 - The second term must minimize the cost on $[\tau, \infty)$, i.e. $\int_\tau^\infty L(\mathbf{x}(t), \phi^*(t)) dt = \min_{\mathbf{u} \in \mathcal{U}} \int_\tau^\infty L(\mathbf{x}(t), \mathbf{u}(t)) dt$, since if this is not the case, then if we decrease τ , then $V(\mathbf{x})$ will decrease, violating our assumption that it is optimal
 - This means that the control law must be optimal at each point in time, so we can consider only a very small initial τ (i.e. the initial step) and ignore the rest
 - This is known as *Bellman's principle of optimality*: an optimal policy has the property that, whatever the initial state and decisions so far are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the decisions so far
 - Therefore $V(\mathbf{x}) = \min_{\mathbf{u} \in \mathcal{U}} \int_0^\tau L(\mathbf{x}(t), \mathbf{u}(t)) dt + V(\mathbf{x}(\tau))$, where $V(\mathbf{x}(\tau))$ is the cost from time τ

onwards and known as the *cost-to-go*

- For small τ , $\int_0^\tau L(\mathbf{x}(t), \mathbf{u}(t)) dt = \tau L(\mathbf{x}, \mathbf{u}) + O(\tau)$ where O has the property that $\lim_{s \rightarrow 0} \frac{O(s)}{s} = 0$
- Using a first-order Taylor series at $\tau = 0$, $V(\mathbf{x}(\tau)) = V(\mathbf{x}) + \frac{d}{dt}V(\mathbf{x}(t))\Big|_{t=0} \tau + O(\tau)$
 - * Note $\frac{d}{dt}V(\mathbf{x}(t)) = \begin{bmatrix} \frac{\partial V}{\partial x_1}(x) & \dots & \frac{\partial V}{\partial x_n}(x) \end{bmatrix} \cdot \frac{d\mathbf{x}}{dt} = \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \cdot (\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u})$
- Substituting both expansions into the expression earlier and some manipulation later we get the HJB equation
- The HJB equation is quadratic in \mathbf{u} , and since \mathbf{R} is positive definite there is a unique minimum, i.e. the problem is convex, with the unique solution given by $\mathbf{u}^* = -\frac{1}{2}\mathbf{R}^{-1}\mathbf{B}^T \left(\frac{\partial V}{\partial \mathbf{x}}(\mathbf{x})\right)^T$
- Propose a trial solution $V(\mathbf{x}) = \mathbf{x}^T \mathbf{P} \mathbf{x}$, where $\mathbf{P}^T = \mathbf{P}$ is positive semidefinite, so $\frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) = 2\mathbf{x}^T \mathbf{P}$ and $\mathbf{u}^* = -\mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}$
 - Substitute into the HJB equation to get $\mathbf{x}^T(-\mathbf{P}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^T\mathbf{P} + \mathbf{P}\mathbf{A} + \mathbf{A}^T\mathbf{P} + \mathbf{Q})\mathbf{x} = 0$
 - Since this holds for any arbitrary $\mathbf{x} \in \mathbb{R}^n$, this means that $-\mathbf{P}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^T\mathbf{P} + \mathbf{P}\mathbf{A} + \mathbf{A}^T\mathbf{P} + \mathbf{Q} = 0$
 - $-\mathbf{P}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^T\mathbf{P} + \mathbf{P}\mathbf{A} + \mathbf{A}^T\mathbf{P} + \mathbf{Q} = 0$ is known as the *algebraic Riccati equation*
- Now we need to show that our candidate optimal control law $\phi(t) = -\mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}(t)$ does lead to the lowest cost, and stabilizes the system so it is admissible
- We will use the result that if \mathbf{Q} is symmetric positive semidefinite, then there exists a matrix square root $\mathbf{Q}^{\frac{1}{2}}$ such that $(\mathbf{Q}^{\frac{1}{2}})^T \mathbf{Q}^{\frac{1}{2}} = \mathbf{Q}$
 - Since \mathbf{Q} is PSD, $\mathbf{Q} = \mathbf{M}^T \mathbf{\Lambda} \mathbf{M}$ where \mathbf{M} is an orthogonal matrix, so $\mathbf{Q}^{\frac{1}{2}} = \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{M}$ where $\mathbf{\Lambda}^{\frac{1}{2}}$ has square roots on the diagonal
- Also, all principal sub-matrices of \mathbf{Q} , i.e. any sub-matrix on the diagonal of \mathbf{Q} , must also be positive semidefinite (this can be shown by noting $\mathbf{x}^T \mathbf{Q} \mathbf{x} \geq 0$ and we can choose \mathbf{x} to isolate any sub-matrix)

Theorem

If (\mathbf{A}, \mathbf{B}) is stabilizable and $(\mathbf{Q}^{\frac{1}{2}}, \mathbf{A})$ is detectable, then the algebraic Riccati equation

$$-\mathbf{P}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^T\mathbf{P} + \mathbf{P}\mathbf{A} + \mathbf{A}^T\mathbf{P} + \mathbf{Q} = 0$$

has a unique symmetric positive semidefinite solution \mathbf{P} , and the control policy

$$\phi^*(t) = -\mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}(t)$$

is the optimal controller for the linear quadratic optimal control problem, and the optimal cost is given by $V(\mathbf{x}) = J(\mathbf{x}, \phi^*) = \mathbf{x}^T \mathbf{P} \mathbf{x}$.

- Proof: Let $\phi \in \mathcal{U}$ and $\mathbf{x}(0) = \mathbf{x}$, we substitute the algebraic Riccati equation for \mathbf{Q} into the cost, then complete the square and use $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\phi(t)$

$$\begin{aligned} -J(\mathbf{x}, \phi) &= \int_0^\infty \mathbf{x}(t)^T \mathbf{Q} \mathbf{x}(t) + \phi(t)^T \mathbf{R} \phi(t) dt \\ &= \int_0^\infty \mathbf{x}(t)^T (\mathbf{P}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^T\mathbf{P} - \mathbf{P}\mathbf{A} - \mathbf{A}^T\mathbf{P}) \mathbf{x}(t) + \phi(t)^T \mathbf{R} \phi(t) dt \\ &= \int_0^\infty (\phi(t) + \mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}(t))^T \mathbf{R} (\phi(t) + \mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}(t)) - (\dot{\mathbf{x}}(t)^T \mathbf{P} \mathbf{x}(t) + \mathbf{x}(t)^T \mathbf{P} \dot{\mathbf{x}}(t)) dt \\ &= \int_0^\infty (\phi(t) + \mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}(t))^T \mathbf{R} (\phi(t) + \mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}(t)) dt - \int_0^\infty \frac{d}{dt} (\mathbf{x}(t)^T \mathbf{P} \mathbf{x}(t)) dt \\ &= \mathbf{x}^T \mathbf{P} \mathbf{x} + \int_0^\infty (\phi(t) + \mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}(t))^T \mathbf{R} (\phi(t) + \mathbf{R}^{-1}\mathbf{B}^T \mathbf{P} \mathbf{x}(t)) dt \end{aligned}$$

- By the positive semidefiniteness of \mathbf{R} , the term inside the integral is always nonnegative, and it is

minimized when $\phi(t) + \mathbf{R}^{-1} \mathbf{B}^T \mathbf{P} \mathbf{x}(t)$, i.e. $\phi^*(t) = -\mathbf{R}^{-1} \mathbf{B}^T \mathbf{P} \mathbf{x}(t)$, which gives us the optimal cost $\mathbf{x}^T \mathbf{P} \mathbf{x}$

Summary

For a system $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$, linear quadratic optimal control seeks to optimize a cost functional

$$J(\mathbf{x}, \phi) = \int_0^\infty \mathbf{x}(t)^T \mathbf{Q} \mathbf{x}(t) + \phi(t)^T \mathbf{R} \phi(t) dt$$

where $\mathbf{u} = \phi$, $\mathbf{x}(0) = \mathbf{x}$, assuming that (\mathbf{A}, \mathbf{B}) is stabilizable and $(\mathbf{Q}^{\frac{1}{2}}, \mathbf{A})$ is detectable. Then the optimal control law is given by

$$\phi^*(t) = -\mathbf{R}^{-1} \mathbf{B}^T \mathbf{P} \mathbf{x}(t)$$

where \mathbf{P} is the unique positive semidefinite solution to the algebraic Riccati equation,

$$-\mathbf{P} \mathbf{B} \mathbf{R}^{-1} \mathbf{B}^T \mathbf{P} + \mathbf{P} \mathbf{A} + \mathbf{A}^T \mathbf{P} + \mathbf{Q} = 0$$