CDS 131 Homework 7: Linear Optimal Control

Winter 2025

Due 2/26 at 11:59 PM

Instructions

This homework is divided into three parts:

1. *Optional Exercises: the exercises are entirely optional but are recommended to be completed before looking at the problems. They consist of easier, more computational questions to help you get a feel for the material.*

2. *Required Problems: the problems are the required component of the homework, and might require more work than the exercises to complete.*

3. *Optional Problems: the optional problems are some additional, recommended problems - some of these might go a little beyond the standard course material.*

All you need to turn in is the solutions to the required problems - the others are recommended but not required.

1 Optional Exercises

1.1 Practice with Dynamic Programming

In this problem, we'll get some practice applying the method of dynamic programming to a simple, scalar optimal control problem. Using the method of dynamic programming, find an optimal state feedback control law which solves the optimal control problem,

$$\inf_{u[0:1]\subseteq\mathbb{R},x[0:2]\subseteq\mathbb{R}} x[2]^2 + \sum_{k=0}^{1} x[k]^2 + u[k]^2, \text{ s.t } x[k+1] = ax[k] + bu[k],\ x[0] = x_0, \tag{1}$$

where $x[k], u[k] \in \mathbb{R}$, and $a, b \in \mathbb{R}$ are fixed constants. Verify that your solution is consistent with that of the general discrete-time LQR problem.

1.2 Practice with the HJB Equation

In this problem, we'll get some practice applying the HJB equation to solve a simple, scalar case of the continuous-time LQR problem. Using the HJB equation, find an optimal state feedback control law which solves the optimal control problem,

$$\inf_{u(\cdot)\in PC(\mathbb{R},\mathbb{R}),x(\cdot)\in C(\mathbb{R},\mathbb{R})} x(T)^2 + \int_0^T x(\tau)^2 + u(\tau)^2 d\tau, \text{ s.t. } \dot{x}(t) = ax(t) + bu(t), x(0) = x_0, \tag{2}$$

where $x(t), u(t) \in \mathbb{R}$ and $a, b \in \mathbb{R}$ are fixed constants. You may leave your answer in terms of the solution to a differential equation if need be. Verify that your solution is consistent with that of the general continuous-time LQR problem.

2 Required Problems

2.1 Linear Quadratic Tracking

In the LQR problem, we seek to regulate the state of the system—that is, keep the state close to the origin. Let's consider a generalization of the LQR problem to the linear quadratic *tracking* problem. Let $N \in \mathbb{Z}_{>0}$, and consider a map $x_d[\cdot] : \mathbb{Z}_{[0,N]} \to \mathbb{R}^n$, representing a desired state trajectory which we would like our system to follow. Our goal is to solve the linear quadratic tracking problem,

$$\inf_{u[0:N-1]\subseteq\mathbb{R}^m, x[0:N]\subseteq\mathbb{R}^n} e[N]^\top Q_f e[N] + \sum_{k=0}^{N-1} e[k]^\top Q e[k] + u[k]^\top R u[k] \tag{3}$$

$$\text{s.t. } x[k+1] = Ax[k] + Bu[k],\ k \in [0, N-1],\ x[0] = x_0. \tag{4}$$

where $e[k] := x[k] - x_d[k]$ is the tracking error at time k. Like the standard LQR problem, we assume that the weight matrices Q_f, Q, and R satisfy $Q_f, Q \in \mathbb{S}^n$, $Q_f, Q \succeq 0$ and $R \in \mathbb{S}^m$, $R \succ 0$. In this problem, we'll show that each step of the dynamic programming recurrence,

$$J_N(x) = (x - x_d[N])^\top Q_f (x - x_d[N]) \tag{5}$$

$$J_k(x) = \inf_{u\in\mathbb{R}^m} (x - x_d[k])^\top Q (x - x_d[k]) + u^\top R u + J_{k+1}(Ax + Bu), \tag{6}$$

has the form $J_k(x) = x^\top P_k x + 2q_k^\top x + r_k$, where $P_k \in \mathbb{S}^n, P_k \succeq 0$, $q_k \in \mathbb{R}^n$, and $r_k \in \mathbb{R}$. Using this form, we'll derive an optimal linear quadratic tracking control law.

1. Let's execute the base case of the dynamic programming algorithm. Show $J_N(x)$ has the form,

$$J_N(x) = x^\top P_N x + 2q_N^\top x + r_N, \tag{7}$$

   for some $P_N \in \mathbb{S}^n, P_N \succeq 0$, $q_N \in \mathbb{R}^n$, and $r_N \in \mathbb{R}$.

2. Assume that J_{k+1} has the desired form, $J_{k+1}(x) = x^\top P_{k+1} x + 2q_{k+1}^\top x + r_{k+1}$. Show that the optimal feedback control law at time step k can be written as $u[k] = K_k x[k] + g_k$, where

$$K_k = -(R + B^\top P_{k+1} B)^{-1} B^\top P_{k+1} A,\ g_k = -(R + B^\top P_{k+1} B)^{-1} B^\top q_{k+1}. \tag{8}$$

   In other words, the optimal control law is a time-varying, affine function of the state.

3. Using backward induction, show that $J_k(x) = x^\top P_k x + 2q_k^\top x + r_k, k \in [0, N]$. State recursive formulas for each of P_k, q_k, and r_k which are valid for $k \in [0, N-1]$. You do not have to simplify your formulas. You may leave them in terms of $K_k, g_k, x_d[k], P_{k+1}, q_{k+1}, r_{k+1}$, and any constant matrices.

2.2 Receding Horizon LQR

Consider the finite-horizon LQR problem,

$$\inf_{u[0:N-1]\subseteq\mathbb{R}^m, x[0:N]\subseteq\mathbb{R}^n} x[N]^\top Q_f x[N] + \sum_{k=0}^{N-1} x[k]^\top Q x[k] + u[k]^\top R u[k] \tag{9}$$

$$\text{s.t. } x[k+1] = Ax[k] + Bu[k],\ x[0] = x_0. \tag{10}$$

Let $u[k] = K_{LQR}[k]x[k],\ k \in [0, N-1]$ denote the optimal control law for this problem. The *receding horizon* control law associated with this optimal control problem is the control law $u[k] = K_{LQR}[0]x[k], k \geq 0$, which is a static state feedback control law with gain equal to the LQR gain at time 0 in the optimization problem.

1. Propose an algorithm to find the minimum horizon N for which the receding horizon control law $u[k] = K_{LQR}[0]x[k]$ renders the origin of $x[k+1] = Ax[k] + Bu[k]$ globally exponentially stable.

2. Conjecture a set of conditions on the system and LQR problem under which your algorithm will be successful. You do not have to prove that your conjecture is correct, but you should support your answer with some basic reasoning.

3. Provide intuitive answers to the following questions. Why might the receding horizon control law fail to stabilize the system if the horizon is too small? Why might a larger horizon lead to stability?

4. State conditions on the system $x[k+1] = Ax[k] + Bu[k]$ and the weight matrices Q and R of the finite-horizon LQR problem for which there exists a choice of $Q_f \in \mathbb{S}^n, Q_f \succeq 0$ for which $K_{LQR}[k] = K_{LQR}[0]$ for all $k \in [0, N-1]$. If you think no such conditions exist, explain why.

2.3 Gain Margin for a Linear Quadratic Regulator

Consider a linear, time-invariant system $\dot{x}(t) = Ax(t) + Bu(t)$. In order for the infinite horizon LQR problem,

$$\inf_{u(\cdot),x(\cdot)} \int_0^\infty x(\tau)^\top Q x(\tau) + u(\tau)^\top R u(\tau) d\tau, \text{ s.t. } \dot{x}(t) = Ax(t) + Bu(t),\ x(0) = x_0, \tag{11}$$

to be feasible, it's sufficient (but not strictly necessary) to have (A, B) controllable and (A, C) is observable, where C is the matrix for which $Q = C^\top C$. Under these conditions, the optimal control law for this problem is given by $u(t) = K_{LQR,\infty}x(t) = -R^{-1}B^\top Px(t)$, where P is the unique, positive definite solution to the *continuous-time algebraic Riccati equation (CARE)*,

$$Q + A^\top P + PA - PBR^{-1}B^\top P = 0. \tag{12}$$

In this problem, we'll prove some basic robustness properties of the infinite-horizon LQR control law using the structure of the continuous-time algebraic Riccati equation.

1. Consider the setting of the infinite-horizon LQR problem posed above. Show that the unique symmetric, positive definite solution P of the continuous-time algebraic Riccati equation is such that

$$(A + BK_{LQR,\infty})^\top P + P(A + BK_{LQR,\infty}) = -PBR^{-1}B^\top P - Q. \tag{13}$$

Using this equality, prove that $A + BK_{LQR,\infty}$ is Hurwitz. *Hints: remember that Q is only PSD—you cannot directly apply the CTLE theorem. How can you use observability of (A, C) to work around this?*

2. Let's derive a simple *gain margin* for the LQR control law—a gain margin considers how much we can amplify our control law before instability. Consider the control law $u = \alpha K_{LQR,\infty}x$, where $\alpha > 0$. Show that for $\alpha > 1/2$, the origin of the closed-loop system $\dot{x}(t) = Ax(t) + Bu(t)$, $u(t) = \alpha K_{LQR,\infty}x(t)$ is globally exponentially stable.

In the language of classical control theory, the ability of a system to remain stable when its control law is scaled by $\alpha \in (1/2, \infty)$ is described as a negative gain margin of 6dB and an infinite positive gain margin.

3 Optional Problems

3.1 Integral Control Design via LQR

In this problem, we'll consider the design of an *integral control law* for a continuous-time system using LQR. Recall that for a system $\dot{x}(t) = Ax(t) + Bu(t)$, the continuous-time LQR problem is defined,

$$\inf_{u(\cdot),x(\cdot)} x(T)^\top Q_f x(T) + \int_0^T x(\tau)^\top Q x(\tau) + u(\tau)^\top R u(\tau) \tag{14}$$

$$\text{s.t. } \dot{x}(t) = Ax(t) + Bu(t),\ x(0) = x_0, \tag{15}$$

where $Q_f, Q \in \mathbb{S}^n, Q_f, Q \succeq 0$ and $R \in \mathbb{S}^m, R \succ 0$. The solution to this problem is given by the control law $u(t) = -R^{-1}B^\top P(t)x(t)$, where $P(t)$ is the solution to the continuous-time Riccati differential equation,

$$\dot{P}(t) = -P(t)A + A^\top P(t) + P(t)BR^{-1}B^\top P(t) - Q,\ P(T) = Q_f \tag{16}$$

Let's consider a slight variation on the standard LQR problem. Define an output $y(t) = Cx(t)$ for the system, where $C \in \mathbb{R}^{p\times n}$, and a signal $z(t) = \int_0^t y(\tau)d\tau$. Consider the optimal control problem,

$$\inf_{u(\cdot),x(\cdot),y(\cdot)} y(T)^\top y(T) + \int_0^T y(\tau)^\top y(\tau) + \rho u(\tau)^\top u(\tau) + \gamma z(\tau)^\top z(\tau) d\tau \tag{17}$$

$$\text{s.t. } \dot{x}(t) = Ax(t) + Bu(t),\ y(t) = Cx(t),\ x(0) = x_0. \tag{18}$$

where ρ and γ are given positive constants. This is similar to an LQR cost function, with the addition of a term that penalizes the norm squared of the integral of y. Compute an optimal control law solving this optimal control problem. Show that the optimal control law can be expressed in the form,

$$u(t) = K_P(t)x(t) + K_I(t)z(t), \tag{19}$$

where $K_P(\cdot) : [0,T] \to \mathbb{R}^{m\times n}$ and $K_I(\cdot) : [0,T] \to \mathbb{R}^{m\times p}$. *Hint: can you recast the given problem as an LQR problem for a different system?*