

Stanford
AA 203: Optimal and Learning-based Control
Problem Set 1, due April 21 by 5:00 pm

Problem 1: Introduction to MRAC

Consider the continuous-time system

$$\dot{y}(t) + \alpha y(t) = \beta u(t).$$

We want to control this system, but we do not know the true plant parameters $\alpha, \beta \in \mathbb{R}$. In this problem, we will use *direct Model-Reference Adaptive Control (MRAC)* to match the behavior of the true plant with that of the reference model

$$\dot{y}_m(t) + \alpha_m y_m(t) = \beta_m r(t)$$

where $\alpha_m, \beta_m \in \mathbb{R}$ are *known* constant parameters, and $r(t)$ is a chosen *bounded* reference signal.

a) Consider the control law

$$u(t) = k_r(t)r(t) + k_y(t)y(t),$$

where $k_r(t)$ and $k_y(t)$ are time-varying feedback gains. Write out the differential equation for the resulting closed-loop dynamics. Use this to verify that, if $y(0) = y_m(0)$ and we knew α and β , the constant control gains

$$k_r^* := \frac{\beta_m}{\beta}, \quad k_y^* := \frac{\alpha - \alpha_m}{\beta}$$

would make the true plant dynamics perfectly match the reference model.

b) When we do not know α and β , we adaptively update our controller over time in response to measurements of $y(t)$. Specifically, we want an *adaptation law* for $k_r(t)$ and $k_y(t)$ to make $y(t)$ tend towards $y_m(t)$ asymptotically. For this, we define the tracking error $e(t) := y(t) - y_m(t)$ and the parameter errors

$$\delta_r(t) := k_r(t) - k_r^*, \quad \delta_y(t) := k_y(t) - k_y^*.$$

Determine a differential equation for e in terms of $e, \dot{e}, y, r, \delta_y, \delta_r$, and suitable constants.

We consider the adaptation law for k_r and k_y described by

$$\begin{aligned} \dot{k}_r(t) &= -\text{sign}(\beta)\gamma e(t)r(t) \\ \dot{k}_y(t) &= -\text{sign}(\beta)\gamma e(t)y(t) \end{aligned},$$

where $\gamma > 0$ is a chosen constant *adaptation gain*. Since we are adapting the gains k_r and k_y of our controller directly, rather than estimates of the system parameters α and β , this is a *direct* adaptation law. We must at least know the sign of β , which indicates in what direction the input $u(t)$ “pushes” the output $y(t)$. For example, when modeling a car, you could reasonably assume that an increased braking force slows down the car. To show that the tracking error and parameter errors are stabilized by our chosen control law and adaptation law, we use Lyapunov theory.

Theorem 1 (Lyapunov). *Consider the continuous-time system $\dot{x} = f(x, t)$, where $x = 0$ is an equilibrium point, i.e., $f(0, t) \equiv 0$. Suppose there exists a continuously differentiable scalar function $V(x, t)$ such that V is positive-definite in x for each $t \geq 0$, and \dot{V} is negative semi-definite in x for each $t \geq 0$. Then $x = 0$ is a stable point in the sense of Lyapunov, i.e., $\|x(t)\|_2$ remains bounded as long as $\|x(0)\|_2$ is bounded.*

c) Consider the state $x := (e, \delta_r, \delta_y)$ and the Lyapunov function candidate

$$V(x) = \frac{1}{2}e^2 + \frac{|\beta|}{2\gamma}(\delta_r^2 + \delta_y^2).$$

Show that $\dot{V} = -\alpha_m e^2$. Based on Lyapunov theory, what can you say about $e(t)$, $\delta_r(t)$, and $\delta_y(t)$ for all $t \geq 0$ if $\alpha_m > 0$?

In general, adaptive controllers yield time-varying closed-loop dynamics, even for LTI systems. As a result, we require more mathematical machinery beyond basic Lyapunov theory to establish anything stronger than Lyapunov stability. To this end, we use Barbalat’s Lemma.

Theorem 2 (Barbalat’s Lemma). *Suppose $g : \mathbb{R} \rightarrow \mathbb{R}$ is differentiable. If g has a finite limit as $t \rightarrow \infty$ and \dot{g} is uniformly continuous, then $\lim_{t \rightarrow \infty} \dot{g}(t) = 0$.*

Boundedness of the derivative of a function is a sufficient condition for Lipschitz continuity and hence uniform continuity. As a result, we have the following corollary.

Corollary 1. *Suppose $g : \mathbb{R} \rightarrow \mathbb{R}$ is twice differentiable. If g has a finite limit as $t \rightarrow \infty$ and \ddot{g} is bounded, then $\lim_{t \rightarrow \infty} \dot{g}(t) = 0$.*

d) Apply Barbalat’s Lemma to V to prove a stronger statement about $e(t)$ than we could originally make with basic Lyapunov theory in part (c).

With the given control law and adaptation law, MRAC proceeds as follows. We choose a reference signal $r(t)$ to excite the reference output $y_m(t)$ and construct the input signal $u(t)$. We use $u(t)$ to excite the true model, from which the output $y(t)$ and tracking error $e(t)$ are observed. The output $y(t)$ is fed back into the control law, while the tracking error $e(t)$ is fed into the adaptation law.

e) Apply MRAC to the unstable plant

$$\dot{y}(t) - y(t) = 3u(t).$$

That is, simulate an adaptive controller for this system that does not have access to the true model parameters $\alpha = -1$ and $\beta = 3$. The desired reference model is

$$\dot{y}_m(t) + 4y_m(t) = 4r(t),$$

with $\alpha_m = 4$ and $\beta_m = 4$. Use an adaptation gain of $\gamma = 2$, and zero initial conditions for y , y_m , k_r , and k_y . For $t \in [0, 10]$, plot both $y(t)$ and $y_m(t)$ in one figure, and $k_r(t)$, k_r^* , $k_y(t)$, and k_y^* in another figure for $r(t) \equiv 4$. Then repeat this for $r(t) = 4\sin(3t)$. Overall, you should have four figures in total. What do you notice about the trends for different reference signals? Why do you think this occurs? In your explanation, try to link your observations with the statements about $e(t)$, $\delta_r(t)$, and $\delta_y(t)$ we were able and *unable* to prove in parts (c,d).

Problem 2: Shortest path through a grid

Consider the shortest path problem in Figure 1 where it is only possible to travel to the right and the numbers represent the travel times for each segment. The control input is the decision to go “up-right” or “down-right” at each node.




Figure 1: Shortest path problem on a grid.

- a) Use Dynamic Programming (DP) to find the shortest path from A to B .
- b) Consider a generalized version of the shortest path problem in Figure 1 where the grid has n segments on each side. Find the number of computations required by an exhaustive search algorithm (i.e., the number of routes that such an algorithm would need to evaluate) and the number of computations required by a DP algorithm (i.e., the number of DP evaluations). For example, for $n = 3$ as in Figure 1, an exhaustive search algorithm requires 20 computations, while the DP algorithm requires only 15.

Problem 3: Machine maintenance

Suppose we have a machine that is either running or is broken down. If it runs throughout one week, it makes a gross profit of \$100. If it fails during the week, gross profit is zero. If it is running at the start of the week and we perform preventive maintenance, the probability that it will fail during the week is 0.4. If we do not perform such maintenance, the probability of failure is 0.7. However, maintenance will cost \$20. When the machine is broken down at the start of the week, it may either be repaired at a cost of \$40, in which case it will fail during the week with a probability of 0.4, or it may be replaced at a cost of \$150 by a new machine that is guaranteed to run through its first week of operation. Find the optimal repair, replacement, and maintenance policy that maximizes total profit over four weeks, assuming a new machine at the start of the first week.

Problem 4: Markovian drone

In this problem, we will apply techniques for solving a Markov Decision Process (MDP) to guide a flying drone to its destination through a storm. The world is represented as an $n \times n$ grid, i.e., the state space is

$$\mathcal{S} := \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1, x_2 \in \{0, 1, \dots, n-1\}\}.$$

In these coordinates, $(0, 0)$ represents the bottom left corner of the map and $(n-1, n-1)$ represents the top right corner of the map. From any location $x = (x_1, x_2) \in \mathcal{S}$, the drone has four possible directions it can move in, i.e.,

$$\mathcal{A} := \{\text{up, down, left, right}\}.$$

The corresponding state changes for each action are:

- **up:** $(x_1, x_2) \mapsto (x_1, x_2 + 1)$
- **down:** $(x_1, x_2) \mapsto (x_1, x_2 - 1)$
- **left:** $(x_1, x_2) \mapsto (x_1 - 1, x_2)$
- **right:** $(x_1, x_2) \mapsto (x_1 + 1, x_2)$

Additionally, there is a storm centered at $x_{\text{eye}} \in \mathcal{S}$. The storm's influence is strongest at its center and decays farther from the center according to the equation $\omega(x) = \exp\left(-\frac{\|x - x_{\text{eye}}\|_2^2}{2\sigma^2}\right)$. Given its current state x and action a , the drone's next state is determined as follows:

- With probability $\omega(x)$, the storm will cause the drone to move in a uniformly random direction.
- With probability $1 - \omega(x)$, the drone will move in the direction specified by the action.
- If the resulting movement would cause the drone to leave \mathcal{S} , then it will not move at all. For example, if the drone is on the right boundary of the map, then moving right will do nothing.

The quadrotor’s objective is to reach $x_{\text{goal}} \in \mathcal{S}$, so the reward function is the indicator function $R(x) = I_{x_{\text{goal}}}(x)$. In other words, the drone will receive a reward of 1 if it reaches the $x_{\text{goal}} \in \mathcal{S}$, and a reward of 0 otherwise. The reward of a trajectory in this infinite horizon problem is a discounted sum of the rewards earned in each timestep, with discount factor $\gamma \in (0, 1)$.

- a) Given $n = 20$, $\sigma = 10$, $\gamma = 0.95$, $x_{\text{eye}} = (15, 15)$, and $x_{\text{goal}} = (19, 9)$, write code that uses value iteration to find the optimal value function for the drone to navigate the storm. Recall that value iteration repeats the Bellman update

$$V(x) \leftarrow \max_{a \in \mathcal{A}} \left(\sum_{x' \in \mathcal{S}} p(x'; x, a) (R(x') + \gamma V(x')) \right)$$

until convergence, where $p(x'; x, a)$ is the probability distribution of the next state being x' after taking action a in state x , and R is the reward function. Plot a heatmap of the optimal value function obtained by value iteration over the grid \mathcal{S} , with $x = (0, 0)$ in the bottom left corner, $x = (n - 1, n - 1)$ in the top right corner, the x_1 -axis along the bottom edge, and the x_2 -axis along the left edge.

- b) Recall that a policy π is a mapping $\pi : \mathcal{S} \rightarrow \mathcal{A}$ where $\pi(x)$ specifies the action to be taken should the drone find itself in state x . An optimal value function V^* induces an optimal policy π^* such that

$$\pi^*(x) \in \arg \max_{a \in \mathcal{A}} \left(\sum_{x' \in \mathcal{S}} p(x'; x, a) (R(x') + \gamma V^*(x')) \right)$$

Use the value function you computed in part (a) to compute an optimal policy. Then, use this policy to simulate the MDP starting from over $N = 100$ time steps starting at $x = (0, 19)$. Plot the policy as a heatmap where the actions `{up, down, left, right}` correspond to the values `{0, 1, 2, 3}`, respectively. Plot the simulated drone trajectory overlaid on the policy heatmap, and briefly describe in words what the policy is doing.

Problem 5: Cart-pole balance

In this problem, we will design a controller to balance an inverted pendulum on a cart, i.e., the classic “cart-pole” benchmark. This system has two degrees of freedom corresponding to the horizontal position x of the cart, and the angle θ of the pendulum (where $\theta = 0$ occurs when the pendulum is hanging straight down). We can apply a force $u \in \mathbb{R}$ to push the cart horizontally, where $u > 0$ corresponds to a force in the positive x -direction. With the state $s := (x, \theta, \dot{x}, \dot{\theta}) \in \mathbb{R}^4$, we can write the continuous-time dynamics of the cart-pole system as

$$\dot{s} = f(s, u) = \begin{bmatrix} \dot{x} \\ \dot{\theta} \\ \frac{m_p \sin \theta (\ell \dot{\theta}^2 + g \cos \theta) + u}{m_c + m_p \sin^2 \theta} \\ - \frac{(m_c + m_p) g \sin \theta + m_p \ell \dot{\theta}^2 \sin \theta \cos \theta + u \cos \theta}{(m_c + m_p \sin^2 \theta) \ell} \end{bmatrix},$$

where m_p is the mass of the pendulum, m_c is the mass of the cart, ℓ is the length of the pendulum, and g is the acceleration due to gravity. We can discretize the continuous-time dynamics using Euler integration with a fixed time step Δt to get the approximate discrete-time dynamics

$$s_{k+1} \approx s_k + \Delta t f(s_k, u_k),$$

where s_k and u_k are the state and control input, respectively, at time $t = k\Delta t$.

- a) Consider the upright state $s^* := (0, \pi, 0, 0)$ with $u^* := 0$, and define $\Delta s_k := s_k - s^*$. Linearizing the approximate discrete-time dynamics $s_{k+1} \approx s_k + \Delta t f(s_k, u_k)$ about (s^*, u^*) yields an approximate LTI system of the form

$$\Delta s_{k+1} \approx A \Delta s_k + B u_k.$$

Express A and B in terms of m_p , m_c , ℓ , g , and Δt . You may use the fact that

$$\frac{\partial f}{\partial s}(s^*, u^*) = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & \frac{m_p g}{m_c} & 0 & 0 \\ 0 & \frac{(m_c + m_p)g}{m_c \ell} & 0 & 0 \end{bmatrix}, \quad \frac{\partial f}{\partial u}(s^*, u^*) = \begin{bmatrix} 0 \\ 0 \\ \frac{1}{m_c} \\ \frac{1}{m_c \ell} \end{bmatrix}.$$

We will design a stabilizing LQR controller for this discrete-time LTI system to solve

$$\begin{aligned} & \underset{\{u_k\}_{k=0}^{\infty}}{\text{minimize}} \quad \sum_{k=0}^{\infty} \left(\frac{1}{2} \Delta s_k^\top Q \Delta s_k + \frac{1}{2} u_k^\top R u_k \right), \\ & \text{subject to } \Delta s_{k+1} = A \Delta s_k + B u_k, \quad \forall k \in \mathbb{N}_{\geq 0} \end{aligned}$$

for fixed $Q, R \succ 0$. Recall that after N iterations of the discrete-time Riccati recursion

$$\begin{aligned} K_k &= -(R + B^\top P_{k+1} B)^{-1} B^\top P_{k+1} A \\ P_k &= Q + A^\top P_{k+1} (A + B K_k) \end{aligned},$$

the cost-to-go matrices $\{P_k\}_{k=0}^N$ and the time-varying feedback gains $\{K_k\}_{k=0}^{N-1}$ describe the optimal LQR controller for a finite-horizon version of the problem above. If (A, B) is stabilizable, then these iterates asymptotically converge to some $P_\infty \succ 0$ and K_∞ , such that $(s_0 - s^*)^\top P_\infty (s_0 - s^*) > 0$ is the finite optimal cost-to-go for any initialization s_0 , and $u_k = K_\infty \Delta s_k$ is the *time-invariant* feedback law for the corresponding optimal LQR controller¹.

- b) Write code to approximate P_∞ and K_∞ for the linearized, discretized cart-pole system by initializing $P_\infty = 0$ and then applying the Riccati recursion until convergence with respect to the maximum element-wise norm condition $\|P_{k+1} - P_k\|_{\max} < 10^{-4}$. Use $m_p = 2$ kg, $m_c = 10$ kg, $\ell = 1$ m, $g = 9.81$ m/s², $\Delta t = 0.1$ s, $Q = I_4$, and $R = I_1$. Report the value of K_∞ up to two decimal places for each entry.

¹The infinite-horizon LQR problem also converges for fixed $Q \succeq 0$ and $R \succ 0$, as long as (A, B) is stabilizable and (A, Q) is observable.

- c) Write code to simulate the continuous-time, nonlinear cart-pole system with the linear feedback controller $u = K_\infty \Delta s$. Initialize the system at $s = (0, 3\pi/4, 0, 0)$, and use a controller sampling rate of 10 Hz. Plot each state variable over time for $t \in [0, 30]$. For your own interest, we provide the function `animate_cartpole` in `animations.py` to create a video animation of the cart-pole over time².

Hint: Write a function `ds = cartpole(s,t,u)` that computes the state derivative `ds` for the continuous-time, nonlinear cart-pole dynamics. To simulate the cart-pole with the fixed control input `u[k]` from state `s[k]` at time `t[k]` to state `s[k+1]` at time `t[k+1]`, you can use the following Python code:

```
from scipy.integrate import odeint
s[k+1] = odeint(cartpole, s[k], t[k:k+2], (u[k],)) [1]
```

Make sure to review the documentation for `odeint`.

- d) To investigate the disturbance rejection ability of the controller, add noise to the system dynamics. Specifically, after each controller sampling period (i.e., every 0.1 s), sample a new noise vector $w \in \mathbb{R}^4$ from the Gaussian distribution with mean $\mu = 0$ and covariance $\Sigma = \text{diag}(0, 0, 10^{-4}, 10^{-4})$, and add it to the state. Simulate the noisy system and plot each state variable over time for $t \in [0, 30]$.

Learning goals for this problem set:

Problem 1: To explore the theoretical underpinnings of MRAC, and observe its behaviour on an example system in simulation.

Problem 2: To familiarize with the DP algorithm and to appreciate the computational savings of DP versus an exhaustive search algorithm.

Problem 3: To apply dynamic programming in stochastic environments by reasoning about expected utilities.

Problem 4: To solve a stochastic optimization problem with value iteration by formulating it as a MDP.

Problem 5: To gain experience with implementing LQR controllers.

²See <https://github.com/StanfordASL/AA203-Homework>.