



# Autonomous Robotics

## Winter 2025

Abhishek Gupta

TAs: Carolina Higuera, Entong Su, Bernie Zhu



# Class Outline

## State Estimation

Robotic System Design

Filtering

Localization

SLAM

## Control

Feedback Control

PID Control

MPC

LQR

## Planning

Search

Heuristic Search

Motion Planning

Lazy Search

## Learning

Imitation Learning

Policy Gradient

Actor-Critic

Model-Based RL

# Logistics

---

- Reading discussions due Wed Feb 12
- Seeded Paper Discussion 2 Monday Feb 17 – emails sent
  
- Post questions, discuss any issues you are having on Ed.
- Students with **no** access to 002, e-mail us with your student ID.
- Students that have not been added to the class, email [abhgupta@cs.washington.edu](mailto:abhgupta@cs.washington.edu) with the subject-line “Waitlisted for CSE478”

Recap

# Different Control Laws

---


- Proportional-integral-derivative (PID) control
- Pure-pursuit control
- Model-predictive control (MPC)
- Linear-quadratic regulator (LQR)
- And many many more!

# PID Intuition

$$u = - \left( \underbrace{K_p e_{ct}}_{\substack{\text{PROPORTIONAL} \\ \text{(PRESENT)}}} + \underbrace{K_i \int e_{ct} dt}_{\substack{\text{INTEGRAL} \\ \text{(PAST)}}} + \underbrace{K_d \dot{e}_{ct}}_{\substack{\text{DERIVATIVE} \\ \text{(FUTURE)}}} \right)$$


- Proportional: minimize the current error!
- Integral: if I'm accumulating error, try harder!
- Derivative: if I'm going to overshoot, slow down!

# Pure Pursuit Controller



- Assume the car is moving with fixed steering angle

# Pure pursuit: Keep chasing lookahead




1. Find a lookahead and compute arc
2. Move along the arc
3. Go to step 1



# Equations of Motion

RECALL



$$\dot{x} = v \cos \theta$$

$$\dot{y} = v \sin \theta$$


$$\dot{\theta} = \omega = \frac{v}{R} = \frac{v \tan \delta}{L}$$

$$\tan \delta = \frac{L}{R} \rightarrow R = \frac{L}{\tan \delta}$$

# Computing the Arc Radius

$$(R_{pp} - b)^2 + a^2 = R_{pp}^2$$

$$R_{pp} = \frac{a^2 + b^2}{2b}$$




# Computing the Arc Radius

$$R_{pp} = \frac{a^2 + b^2}{2b}$$


$$\begin{bmatrix} a \\ b \end{bmatrix} = R(-\theta) \left( \begin{bmatrix} x_{\text{ref}} \\ y_{\text{ref}} \end{bmatrix} - \begin{bmatrix} x \\ y \end{bmatrix} \right)$$

**Different than cross-track error  
(this is ref. position in robot frame;  
vice versa for cross-track error)**



# Question: How do I choose L?

---



# Controller Design Decisions

---

1. Get a reference path/trajectory to track
2. Pick a reference state from the reference path/trajectory
3. Compute error to reference state
4. Compute control law to minimize error



Option 1:

Bang-bang control



Option 2:

PID control



Option 3:

Pure-pursuit control

Are we done?

# Lecture Outline

---

**Recap of Pure Pursuit**




From tracking to optimal control



Linear Quadratic Regulator

# Controller Design Decisions

---



# Control as an Optimization Problem

---

- For a sequence of  $H$  control actions
  1. Use model to predict consequence of actions (i.e.,  $H$  future states)
  2. Evaluate the cost function
- Compute optimal sequence of  $H$  control actions (minimizes cost)



# Generalized Problem: Optimal Control

- Minimize sum of costs, subject to dynamics and other constraints

$$\begin{aligned} \min_{u_{1:T}} \sum_{t=1}^T c(x_t, u_t) \\ \text{s.t. } x_{t+1} = f(x_t, u_t) \end{aligned}$$

Can be costs like smoothness, preferences, speed

Can be constraints like velocity/acceleration bounds

# Linear Quadratic Regulator

---

- **Linear** system (model)
- **Quadratic** cost function to minimize

$$x_{t+1} = Ax_t + Bu_t$$
$$\sum_t x_t^\top Q x_t + u_t^\top R u_t$$

# Linear System

- **Linear** system (model)
- **Quadratic** cost function to minimize

$$x_{t+1} = Ax_t + Bu_t$$
$$\sum_t x_t^\top Q x_t + u_t^\top R u_t$$


$$x_{t+1} = A x_t + B u_t$$

(N x 1)      (N x N)(N x 1)      (N x M)(M x 1)

**STATE** → **NEXT STATE**

**CONTROL** → **NEXT STATE**

# Example: Inverted Pendulum (Linear System)



$$mgl \sin \theta + \tau = ml^2 \ddot{\theta}$$

$$\ddot{\theta} = \frac{g}{l} \sin \theta + \frac{\tau}{ml^2} \approx \frac{g}{l} \theta + \frac{\tau}{ml^2}$$

$$\begin{bmatrix} \theta_{t+1} \\ \dot{\theta}_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & \Delta t \\ \frac{g}{l} \Delta t & 1 \end{bmatrix} \begin{bmatrix} \theta_t \\ \dot{\theta}_t \end{bmatrix} + \begin{bmatrix} 0 \\ \Delta t \end{bmatrix} \frac{\tau}{ml^2}$$

$x_{t+1} \qquad A \qquad x_t \qquad B \qquad u_t$

# Quadratic Cost Function

- **Linear** system (model)
- **Quadratic** cost function to minimize

$$x_{t+1} = Ax_t + Bu_t$$
$$\sum_t x_t^\top Q x_t + u_t^\top R u_t$$

$$x_t^\top Q x_t$$

$$(1 \times N)(N \times N)(N \times 1)$$


**STATE COST**

$$u_t^\top R u_t$$

$$(1 \times M)(M \times M)(M \times 1)$$

**CONTROL COST**

# Example: Inverted Pendulum (State Cost)




$$x_t^\top Q x_t \quad (\text{QUADRATIC FORM})$$

$$= \begin{bmatrix} \theta_t \\ \dot{\theta}_t \end{bmatrix}^\top \begin{bmatrix} Q_{\theta\theta} & Q_{\theta\dot{\theta}} \\ Q_{\dot{\theta}\theta} & Q_{\dot{\theta}\dot{\theta}} \end{bmatrix} \begin{bmatrix} \theta_t \\ \dot{\theta}_t \end{bmatrix}$$

$$= Q_{\theta\theta}\theta_t^2 + 2Q_{\theta\dot{\theta}}\theta_t\dot{\theta}_t + Q_{\dot{\theta}\dot{\theta}}\dot{\theta}_t^2$$

$$Q \succ 0 \Leftrightarrow z^\top Q z > 0, \forall z \neq 0$$

# Example: Inverted Pendulum (Control Cost)



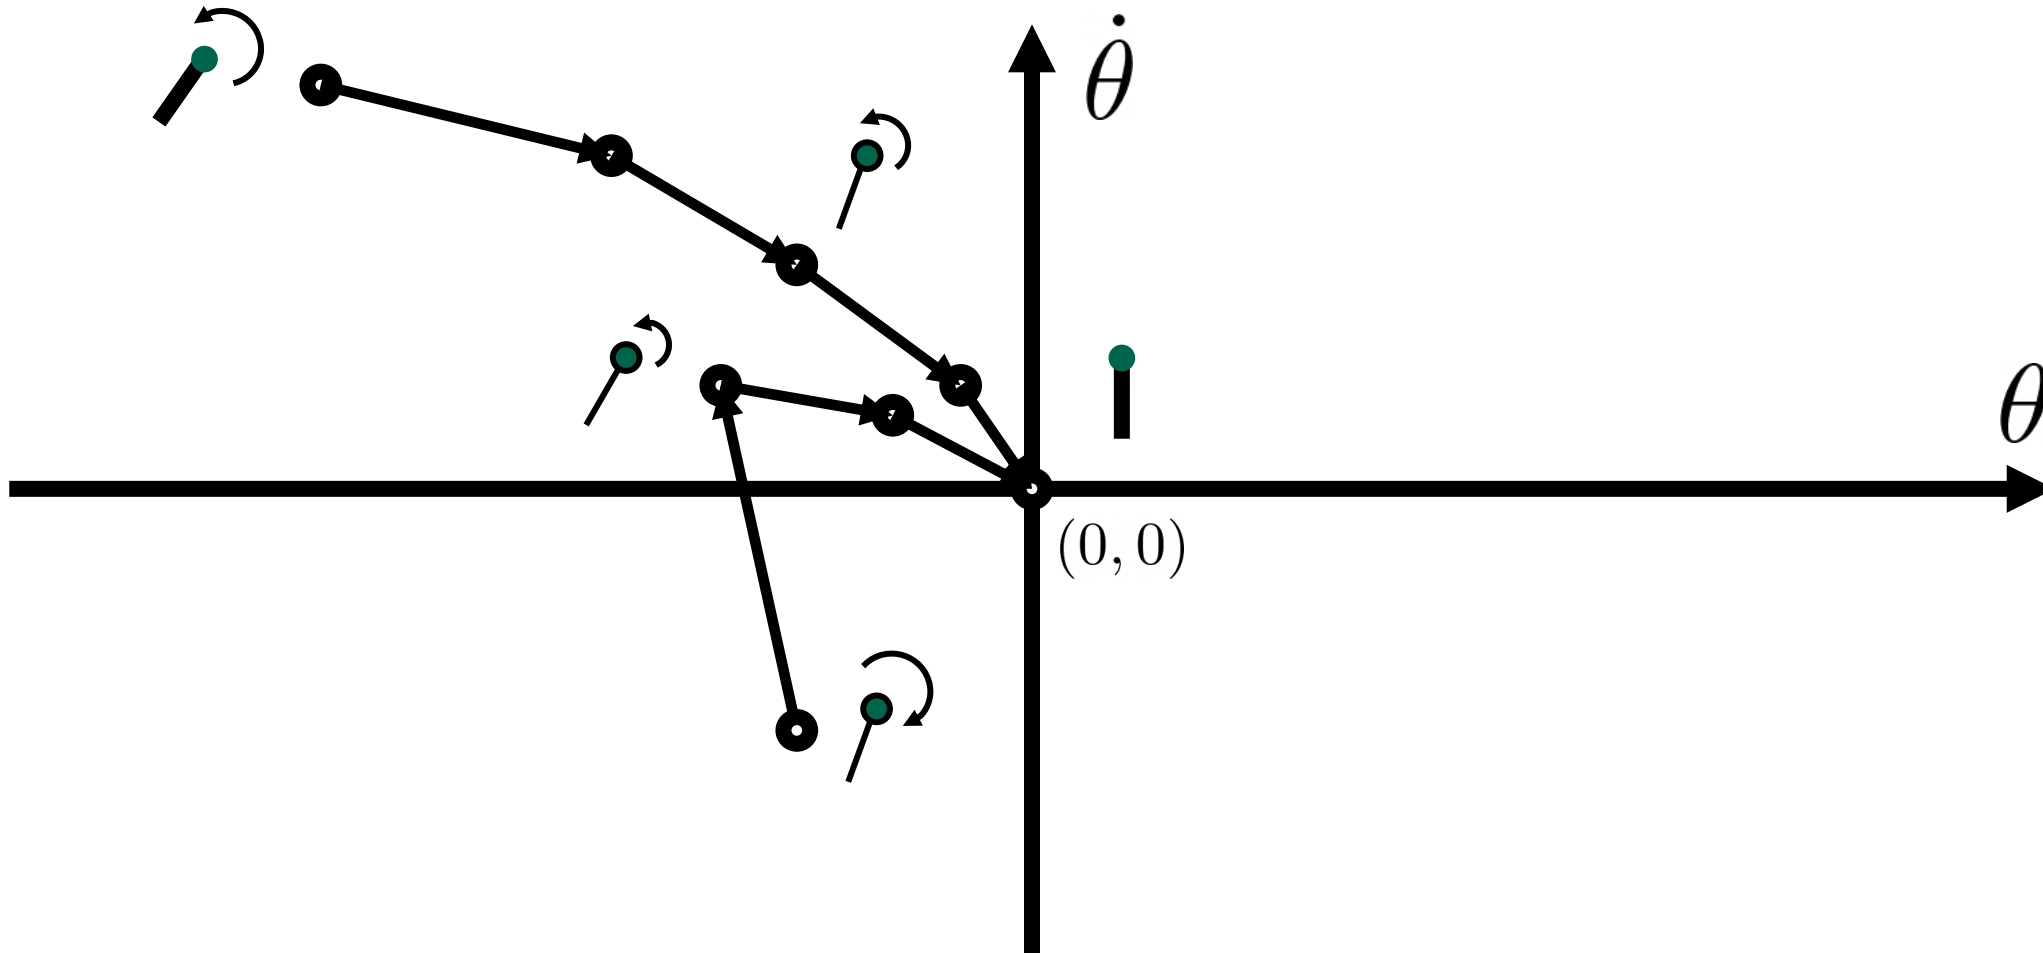
$$u_t^\top R u_t \quad (\text{QUADRATIC FORM})$$

$$= \frac{\tau_t}{ml^2} [R_{\tau\tau}] \frac{\tau_t}{ml^2}$$

$$= R_{\tau\tau} \left( \frac{\tau_t}{ml^2} \right)^2$$

$$R \succ 0 \Leftrightarrow z^\top R z > 0, \quad \forall z \neq 0$$

# Example: Inverted Pendulum






# How do we solve for controls?

Dynamic programming to the rescue!

Start from timestep  $T-1$  and solve backwards



# Bellman Equation for Dynamic Programming

- **Linear** system (model)
- **Quadratic** cost function to minimize

$$x_{t+1} = Ax_t + Bu_t$$
$$\sum_t x_t^\top Q x_t + u_t^\top R u_t$$

$$J^*(x_t) = \min_{u_t} x_t^\top Q x_t + u_t^\top R u_t + J^*(x_{t+1})$$

**MINIMUM COST,  
STARTING FROM**

$x_t$

**IMMEDIATE  
COST**

**MINIMUM FUTURE  
COST, STARTING**

**FROM**  $x_{t+1}$

# Start from the back: Time-to-go = 0

---

$$J_0(x) = \min_u x^\top Q x + u^\top R u$$

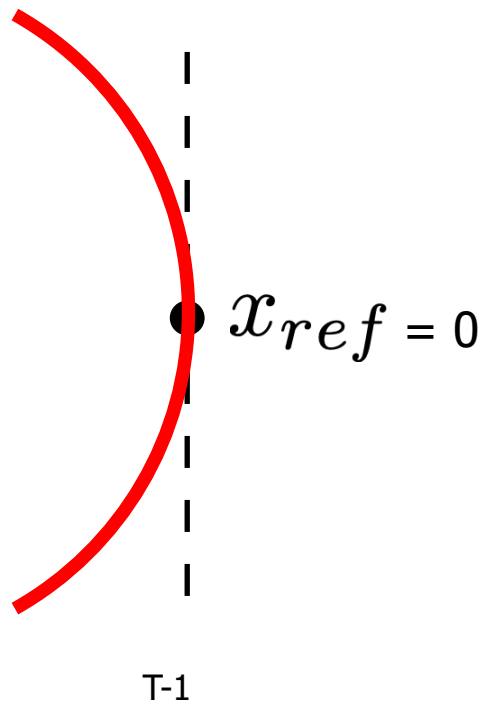
(whiteboard)

# Start from the back: Time-to-go = 0

$$J_0(x) = \min_u x^\top Qx + u^\top Ru = x^\top Qx = x^\top P_0x$$

Minimized with  $u = 0$

$P_0 = Q$




Note that the cost is quadratic in  $x$

# Take one step towards the start: Time-to-go = 1

$$J_0(x) = \min_u x^\top Qx + u^\top Ru = x^\top Qx = x^\top P_0x$$

$$J_1(x) = \min_u x^\top Qx + u^\top Ru + J_0(Ax + Bu)$$



Solve for control at timestep T-1, accounting for impact on the future, through dynamics

# Take one step towards the start: Time-to-go = 1

---

$$J_1(x) = \min_u x^\top Q x + u^\top R u + J_0(Ax + Bu)$$

(Move to whiteboard)

# Value Iteration (Horizon = 1)

---

$$J_1(x) = \min_u [x^\top Qx + u^\top Ru + (Ax + Bu)^\top P_0(Ax + Bu)]$$

$$\nabla_u[\cdot] = 2Ru + 2B^\top P_0(Ax + Bu) = 0$$

$$u = -(R + B^\top P_0 B)^{-1} B^\top P_0 Ax$$

$$J_1(x) = x^\top P_1 x$$

$$P_1 = Q + K_1^\top R K_1 + (A + B K_1)^\top P_0 (A + B K_1)$$

$$K_1 = -(R + B^\top P_0 B)^{-1} B^\top P_0 A$$

# Turns into a recursion at time-to-go = $i$

---

$$K_i = -(R + B^\top P_{i-1} B)^{-1} B^\top P_{i-1} A$$

$$P_i = Q + K_i^\top R K_i + (A + B K_i)^\top P_{i-1} (A + B K_i)$$

$$u = K_i x, \quad J_i(x) = x^\top P_i x$$

Optimal controller is linear in  $x$

Optimal cost is quadratic in  $x$

**RUNTIME:**  $O(H(n^3 + m^3))$



# The LQR algorithm

Algorithm OptimalValueControl(A, B, Q, R, time-to-go):

if time-to-go == 0:

return 0, Q

else:

$P_{i-1} = \text{OptimalValueControl}(A, B, Q, R, \text{time-to-go} - 1)$

$K_i = -(R + B^\top P_{i-1} B)^{-1} B^\top P_{i-1} A$

$P_i = Q + K_i^\top R K_i + (A + B K_i)^\top P_{i-1} (A + B K_i)$

return  $K_i, P_i$

Optimal controller is linear in  $x$

Optimal cost is quadratic in  $x$

# Unpacking LQR intuitively

---

$$K_i = -(R + B^\top P_{i-1} B)^{-1} B^\top P_{i-1} A$$

$$P_i = Q + K_i^\top R K_i + (A + B K_i)^\top P_{i-1} (A + B K_i)$$

$$u = K_i x, \quad J_i(x) = x^\top P_i x$$

# Unpacking LQR intuitively

$$K_i = -(R + B^\top P_{i-1} B)^{-1} B^\top P_{i-1} A$$

Recall Kalman Filtering

$$\frac{B^\top P_{i-1} A}{R + B^\top P_{i-1} B}$$

Set A, B = I

$$\frac{P_{i-1}}{R + P_{i-1}}$$

Tradeoff between future cost  $P_{i-1}$  and current cost R

# Unpacking LQR intuitively

$$x^T \left[ P_i = \underbrace{Q}_{\text{Current state cost}} + \underbrace{K_i^T R K_i}_{\text{Current action cost}} + \underbrace{(A + BK_i)^T P_{i-1} (A + BK_i)}_{\text{Optimal cost in the future based on dynamics}} \right] x$$

The diagram illustrates the decomposition of the LQR cost function. The equation is presented as  $x^T [ \dots ] x$ . Three arrows point from the terms inside the brackets to their respective interpretations:  $Q$  is labeled 'Current state cost',  $K_i^T R K_i$  is labeled 'Current action cost', and  $(A + BK_i)^T P_{i-1} (A + BK_i)$  is labeled 'Optimal cost in the future based on dynamics'.

# Linear Quadratic Regulator

---

- For **linear** systems with **quadratic** costs, we can write down very efficient algorithms that return the optimal sequence of actions!
  - Special case where dynamic programming can be applied to continuous states and actions (typically only discrete states and actions)
- Many LQR extensions: non-linear systems, linear time-varying systems, trajectory following for non-linear systems, arbitrary costs, etc.

# LQR in Action: Stanford Helicopter


---



# LQR in Action



Overcoming challenging indoor environments.



# Class Outline

## State Estimation

Robotic System Design

Filtering

Localization

SLAM

## Control

Feedback Control

PID Control

MPC

LQR

## Planning

Search

Heuristic Search

Motion Planning

Lazy Search

## Learning

Imitation Learning

Policy Gradient

Actor-Critic

Model-Based RL