Robotics and Robot Vision

Guide to AI and Data Courses · Part III: Signals, Images and Audio

Kinematics | Control | Localization | Robot vision | Path planning | Neighboring courses

Sercan KÜLCÜ · Giresun University · Computer Engineering

© Sercan Külcü. All rights reserved. · Interactive HTML version of these slides: https://sercankulcu.github.io/files/ai_courses/slides/Chapter_23_Robotics.html

INTRODUCTION

What Question Does It Ask?

> **How does a machine sense the physical world, know where it is and move?**

Robotics studies how to design and control **machines that sense and move** in the physical world. The course combines perception, localization, planning and control under noisy measurements and real-time constraints; robot vision is the eye of this loop.

At a glance

APPROACH sense-plan-act loop	**INPUT** sensor measurements	**OUTPUT** motor commands, path
SUCCESS MEASURE task success, safety	**MATH LOAD** high: transformations, control	**DATA NEEDS** physical experiments and simulation

A Short History

1961
Unimate, the first industrial robot, starts work at a General Motors plant.

1966
The Shakey project starts at SRI: the first mobile robot to sense and plan; A* is developed for it.

1986
Brooks proposes the subsumption architecture: direct sense-act behaviors instead of planning.

1997
The Sojourner rover explores the surface of Mars on the Mars Pathfinder mission.

2002
The Roomba vacuum goes on sale; mobile robots enter homes.

2005
Stanford's Stanley completes and wins the DARPA Grand Challenge desert race with no driver.

THE COURSE

Core Concepts

Degrees of freedom
The number of axes along which a robot can move independently.

Coordinate frame
World, robot, camera; transformation matrices between them.

Forward kinematics
Finding the end-effector position from joint angles.

Inverse kinematics
Finding the joint angles for a desired end-effector position.

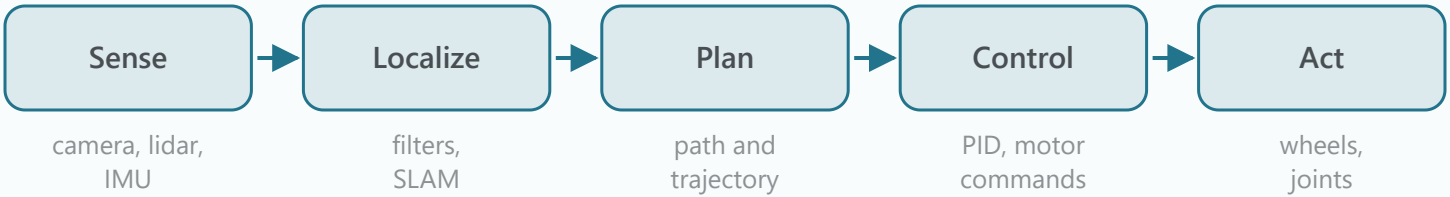
PID control
Feedback that corrects error by proportional, integral and derivative terms.

Localization
Position from noisy measurements with Kalman and particle filters.

SLAM
Building a map while locating yourself in it.

Path planning
A*, RRT: a collision-free path to the goal.

A Typical Workflow



Example: A warehouse robot scans its surroundings with lidar and a camera, estimates its position on the map with a particle filter, plans a path between shelves with A* and sets wheel speeds with PID. The loop runs tens of times a second; each pass, measurements correct the plan.

THE COURSE

What Does the Course Cover?

1. Position and orientation: rotation matrices, frames
2. Forward kinematics and Denavit-Hartenberg notation
3. Inverse kinematics and the Jacobian
4. Robot dynamics: forces, torques, equations of motion
5. Feedback control: PID and stability
6. Trajectory generation and manipulator control
7. Mobile robot kinematics and odometry
8. Sensors: cameras, lidar, IMU, encoders
9. Probabilistic state estimation: the Kalman filter
10. Particle filters and Monte Carlo localization
11. Mapping and SLAM
12. Motion planning: A*, probabilistic roadmaps, RRT

Skills you will gain

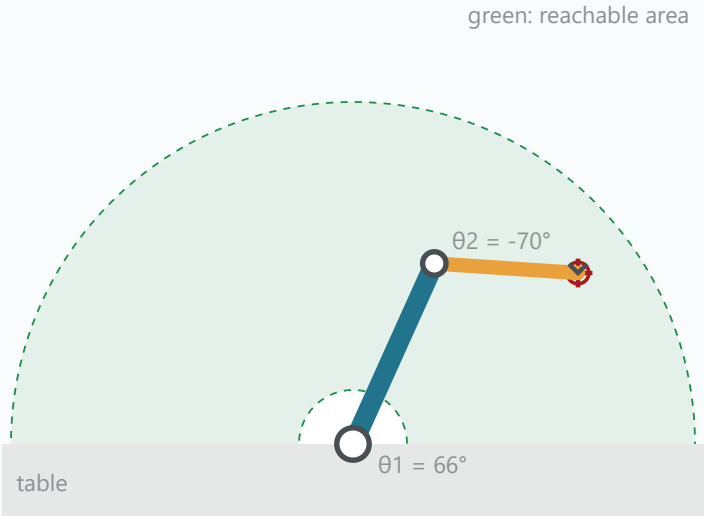
writing ROS 2 nodes and topics | simulating robots in Gazebo | planning arm motion with MoveIt | tuning PID gains | sensor fusion with a Kalman filter | robot vision with OpenCV

A typical semester; the order and coverage vary between universities.

TRY IT

A Robot Arm Picks Up a Part

The tip of a two-joint arm on an assembly line must reach a part on the table, but motors take joint angles, not tip positions. Click a target: inverse kinematics computes the angles and the arm turns there.



Upper arm L1 = **110** mm

Forearm L2 = **80** mm

Elbow up | Elbow down

Target (125, 95) mm, distance from base 157 mm
θ1 = **65.8°**, θ2 = **-69.5°**
Forward kinematics check: tip is 0.0 mm from the target.

Two angle pairs reach the same tip position: elbow up and elbow down. The planner picks the collision-free one.

Robotics turns “where?” into “which angles?”: **inverse kinematics**, workspace limits and choosing the collision-free one among several solutions. A real arm has 6 joints and obstacles.

CONTEXT

Prerequisites

Courses worth knowing before this one, and why they are needed.

Course	Why it is needed
Linear Algebra	coordinate transforms, kinematics
Physics (mechanics)	force, torque, velocity and acceleration: the language of robot dynamics
Microprocessors	reading sensors, driving motors and real-time software
Programming with C	embedded control software is mostly written in C/C++

CONTEXT

Are You Ready for This Course?

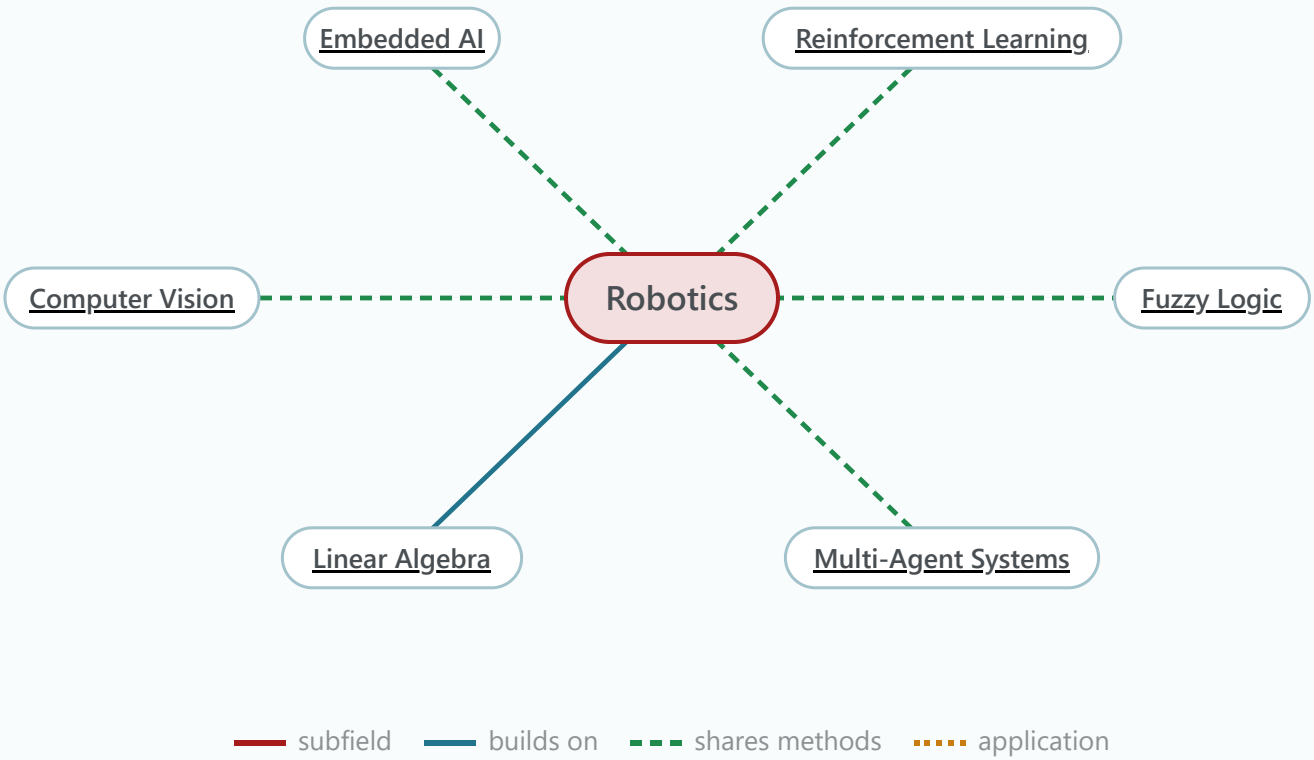
Ask yourself; for every "no", the right-hand column shows where to start.

Ask yourself	If not, start with
Can you move a 2D point into another frame with a rotation matrix?	Linear Algebra
Can you explain velocity, acceleration and torque with an example?	Mechanics from physics: Newton's laws and rotational motion
Can you explain what the mean and variance of a Gaussian mean?	Probability and Statistics
Can you write a class in C++ or Python and process sensor data in a loop?	An intro C course, then a short introduction to classes in Python

If you can say "yes" to most of these, you are ready; one or two gaps can be closed during the semester.

CONTEXT

Neighbouring Courses



OVERLAP

Overlapping Courses

Relation	Course	Shared ground
builds on	This course builds on Linear Algebra	coordinate transforms, kinematics
shares methods	Shares methods with Reinforcement Learning	control, learning motor skills
shares methods	Shares methods with Fuzzy Logic and Fuzzy Systems	fuzzy controllers
shares methods	Shares methods with Multi-Agent Systems	swarm robotics, coordination
shares methods	Shares methods with Computer Vision	robot vision, SLAM
shares methods	Shares methods with Embedded AI	real-time on-device inference

Where It Differs (1/2)

Criterion	Robotics	Computer Vision	Artificial Intelligence
Core question	How does a machine sense the physical world, know where it is and move?	How does a computer understand what is in an image or a video?	How does a machine act intelligently by searching, reasoning and planning?
Approach	sense-plan-act loop	extracting meaning from images	search, logic, probability and learning combined
Input	sensor measurements	image or video	problem definition: states, actions, goal
Output	motor commands, path	label, box, depth, 3D structure	action sequence, plan or decision
Success measure	task success, safety	accuracy, mAP, IoU	reaching the goal, cost, rationality
Math load	high: transformations, control	high: geometry, linear algebra	medium: logic, probability, graphs
Data needs	physical experiments and simulation	many labeled images	little; often a model and knowledge suffice

- **Computer Vision:** Computer vision understands what is in a scene; robot vision turns the image into position and motion decisions.
- **Artificial Intelligence:** AI plans over abstract states; robotics carries out the plan in a continuous, noisy, physical world.

DIFFERENCES

Where It Differs (2/2)

Criterion	Robotics	Reinforcement Learning	Embedded AI
Core question	How does a machine sense the physical world, know where it is and move?	How can good decisions be learned by trial and error from a reward signal alone?	How do AI models run on small, low-power devices?
Approach	sense-plan-act loop	trial and error with rewards	resource-constrained inference
Input	sensor measurements	environment: state, reward	trained model, device constraints
Output	motor commands, path	policy: state → action	small, fast model
Success measure	task success, safety	total expected reward	latency, energy, memory
Math load	high: transformations, control	high: probability, dynamic programming	medium
Data needs	physical experiments and simulation	generated through interaction	training data already exists

- **Reinforcement Learning:** Reinforcement learning learns control by trial and error; robotics mostly designs it from models and control theory.
- **Embedded AI:** Embedded AI fits the model onto a small device; robotics ties that device into a loop with sensors and motors.

PRACTICE

Applications

Industrial arms
Welding, painting, assembly lines.

Warehouse robots
Moving shelves and picking orders.

Autonomous vehicles
Perception, localization and planning.

Surgical robots
Making the surgeon's motion more precise.

Planetary exploration
Mars rovers, autonomous driving at a distance.

Home robots
Robot vacuums and lawn mowers.

PRACTICE

Risks and Ethics

Collisions with people

Robot arms exert large forces; speed and force limits, emergency stops and safety zones are part of collaborative robot design.

Sensor failure and spoofing

A dirty lidar, a camera blind spot or a spoofed GPS signal can lead a robot confidently to the wrong place.

Liability in autonomous crashes

In crashes such as the 2018 Uber test vehicle that killed a pedestrian in Arizona, blame is split between maker, operator and driver.

Surveillance and weapons

Camera-equipped mobile robots and drones open private spaces to monitoring; autonomous weapons are a separate international debate.

In depth: Explainable AI and AI Ethics.

PRACTICE

Tools and Datasets

Tools and libraries

ROS 2 · Gazebo · MoveIt · OpenCV · Point Cloud Library · Webots

Well-known datasets

KITTI · TUM RGB-D · EuRoC MAV · nuScenes · YCB object set

In robotics most experiments run first in **simulation**; moving to a real robot brings friction, noise and latency.

PRACTICE

Common Misconceptions

Common belief	In fact
Robotics is only mechanics and electronics.	Today's hard problems are mostly in software: perception, localization and planning.
A good plan gets the robot to the goal.	Wheels slip and sensors are noisy; a plan holds only with continuous feedback.
What works in simulation works in reality.	The sim-to-real gap is a research topic of its own in robotics.
Robot vision is the same as computer vision.	Robot vision stresses depth, pose and real time; its output is a motion.

NEXT STEPS

After This Course

Courses that build on or extend this one

This course is mostly an end point: it builds on the others.

Related careers

Robotics software engineer | Control engineer | Perception engineer | Autonomous driving engineer | Mechatronics engineer

Robotics brings vision, learning and control together in one physical system. From here you can move to learning motion with reinforcement learning, to multi-agent systems with robot swarms, or to on-device inference with embedded AI.

NEXT STEPS

Further Reading

- Siegwart, Nourbakhsh and Scaramuzza, *Introduction to Autonomous Mobile Robots*, 2nd ed., MIT Press, 2011.
- Thrun, Burgard and Fox, *Probabilistic Robotics*, MIT Press, 2005.
- Lynch and Park, *Modern Robotics: Mechanics, Planning, and Control*, Cambridge University Press, 2017.
- Corke, *Robotics, Vision and Control*, 2nd ed., Springer, 2017.

The End

Chapter 23: Robotics and Robot Vision

← Previous chapter · Map of the field · Next chapter →