Syllabus

Books

1. CORKE - Robotics, Vision and Control: Fundamental Algorithms in PYTHON, by Peter Corke, 3rd edition, 2023. This book is required but is not free to download. See also this repo for the code.
2. LYNCH - Modern Robotics: Mechanics, Planning and Control.
3. THRUN - Probabilistic Robotics, by Sebastian Thrun, Wolfram Burgard, and Dieter Fox, 2005.
4. AIMA - Artificial Intelligence: A Modern Approach, by Stuart Russell, 4th edition, 2021 and also here..

Students after completing this course will be able to

1. To design the various subsystems involved in robotic agents with egomotion.

2. To design and implement perception using sensor fusion such as computer vision together with other sensing streams (such as LIRAR) available to the robot designer.
3. To implement planning algorithms for long-term tasks such as path planning or short term tasks such as motion or trajectory planning.
4. To teach robots a control policy using just a world model simulator and allow the simulated policy to be transferred to the real world.
5. To instruct robots with natural language.
6. To enable students to program all of the above in an industry-standard setting focused around the ROS2 framework.

Planned Schedule

- Part I: Robotic Perception (Lectures 1-5)
- Part II: Localization and Mapping (Lectures 6-7)
- Part III: Motion Planning & Control (Lectures 8-9)
- Part IV: Reinforcement, Imitation Learning and Sim2Real (Lectures 10-12)

Week	Content
Lecture 1	We start with an introduction to AI and robotics from a systems perspective with emphasis on the autonomous vehicles application domain. We explain the various systems that need to be engineered to allow safe and efficient self-driving and review prerequisites on programming (Python) as well as linear algebra and probability theory. With the help of the TAs & other tutorial videos, we also ensure that students have set up a programming environment necessary for the projects and assignments of the course. Reading: course website pages on dev environment setup, probability theory, and linear algebra.
Lecture 2	The perception subsystem is the first stage of the processing chain in robotics. It processes and fuses multi-modal sensory inputs and is implemented using deep neural networks. Our focus here is to understand how prediction can be engineered by taking the maximum likelihood optimization principle and its associated cross-entropy loss function and applying it to neural networks - in this lecture focusing on fully connected and subsequently convolutional architectures for supervised regression and classification tasks. Reading: AIMA Chapter 19 and 21.
Lecture 3	Naturally, the two tasks of classification and regression will be combined to form object detectors. Initially, we cover YOLOv8 and MaskRCNN as representatives of single and two-stage architectures and introduce the engineering aspects of building object detection models with pre-trained feature extractors and using transfer learning for making predictions outside the scope of the pretrained models - a typical setting in robotic applications. Reading: AIMA Chapter 25 (in part).
Lecture 4	In this lecture, we expand on object detection and provide the agent with the additional capabilities of semantic and instance segmentation that are essential

Week	Content
	for completing planning tasks in complex scenes. Reading: AIMA Chapter 25.
Lecture 5	This lecture introduces probabilistic models that process the perceptive predictions over time and understand how the agent will track/update its time-varying belief about the state of the environment. This is achieved with recursive state estimation algorithms acting on a problem setting represented with dynamic Bayesian networks. This lecture introduces Bayesian filters in discrete and continuous state spaces (Kalman filters). All robots employ such filters. Reading: AIMA Chapters 12, 13 & 14.
Lecture 6	In Part I, we built a static agent that can track moving objects - we expand on agents that move themselves (egomotion). We present well-established probabilistic motion models such as the velocity and odometry model and introduce localization as the problem of estimating the agent's pose in the environment. We evaluate various algorithms that can solve such an estimation problem. Reading: CORKE Chapter 4, THRUN Chapter 5 and 7.
Lecture 7	Up to this point, the mobile agent faced a localization estimation task assuming the knowledge of the environment's map. We now expand on agents that can do Simultaneous Localization and Mapping (SLAM). Due to the plethora of the SLAM methods, we focus here on one that brings together Part I and Part II: Visual SLAM demonstrating how even monocular cameras can be used to move the agent safely in a dynamic environment. Reading: CORKE Chapter 6, THRUN Chapters 9 & 10.
Lecture 8	After the successive enhancements of Part I and Part II, the agent now has a clear view of the environment state such as what and where the objects that surround it are, is able to track them as they move, and knows its location as it moves. In this lecture, we introduce the agent's short and longer-term goals and how optimal planning under uncertainty solutions produces the best sequence of actions to reach its goal state. Although our final destination is optimal action-taking in dynamic stochastic environments, we start by connecting planning and problem-solving with the assumption that the goal state is feasible and the environment is known and deterministic. In this case, the planning problem becomes that of a search and covers a range of search-based algorithms such as A*, D*, RRT*, and PRM. Readings: CORKE Chapter 2, AIMA Chapters 3 & 4 and 11.
Lecture 9	This lecture introduces the control subsystem responsible for executing the planned actions. We start with the basics of control theory and introduce the PID controller, then move to more advanced control algorithms such as LQR and MPC. Reading: Course notes.
Lecture 10	We now extend our assumptions: the utility of the agent now depends on a sequence of decisions, and the stochastic environment offers a feedback signal to the agent called reward. We review how the agent's policy, the sequence of

Week	Content
	actions, can be calculated when it fully observes its current state (MDP) and also when it can only partially do so (POMDP). The discussion here is on fundamentals and goes through an understanding of the Bellman expectation backup and optimality equations following David Silver's (DeepMind) lecture series at Oxford. Readings: AIMA Chapter 16 & 17.
Lecture 11	The prediction and control algorithms that approximate the optimal MDP policies using either a world model or a model-free approach are known as Reinforcement Learning (RL). This lecture establishes the connection between MDP and RL and dives into deep learning, focusing on Deep RL algorithms applicable in robotics - mostly model-free methods. Reading: AIMA Chapter 22.
Lecture 12	Simulation-to-reality transfer has emerged as a popular and highly successful method to train robotic control policies for a wide variety of tasks. In this concluding topic, we review various methods proposed to bridge the gap between simulation and reality, such as domain adaptation, and demonstrate a real-world use case based on PyBullet - a popular sim2real tool. Reading: Course notes.
Lecture 13	Review - last lecture before the final.