COL864: Special Topics in Artificial Intelligence
Planning and Estimation for Autonomous Systems
Credits: (3-0-0)
Holi Term 2021

Description

Planning and estimation are central to modern autonomous systems. This course will cover the concepts, principles and methods for intelligent decision-making with imperfect or uncertain knowledge. Students will develop an understanding of how different planning and learning techniques are useful in problem domains where robots or other embodied-AI agents are deployed. Previous coursework in artificial intelligence or machine learning is required.

Topic list (tentative)

- Introduction
  AI view of autonomous systems: Centrality of estimation and planning.

- State Estimation
  Graphical models (review), Bayes Filter and Kalman Filter

- Task Planning
  STRIPS-planning, PDDL planning, Graph Plan

- Planning under Uncertainty
  MDPs (review), Tree Search for MDPs

- Reinforcement learning for Robot control
  Monte-carlo planning, Deep RL applications, Imitation Learning

- POMDPs
  Belief states and Policy trees

- Information Gathering and Exploration
  Gaussian Processes and exploration algorithms.

- Others (if time permits)
  Human-robot interaction, applications of neural models, scene understanding etc.

Course Components

Minor and major exams. Programming assignments (tentatively 1-2). Study of a contemporary works in planning and learning technique relevant to autonomous systems (details in due course).

Pre-requisites

Introduction to Artificial Intelligence (COL333-671) or Introduction to Machine Learning (COL774 or equivalent). Programming proficiency and knowledge of probabilistic models, basic deep learning, basic search algorithms, logic and probability will be an advantage.

Learning outcomes

At the end of the course students will model a robotic system (e.g., a ground robot or manipulator) as a decision-making AI agent. Students will be able to formulate/solve relevant planning and estimation problems in this domain and understand how incorporate recent learning-based methods decision-making algorithms.

Other Information

This course will focus on AI aspects of autonomous systems. A robotic system (ground/air vehicle or manipulator) will be modeled as an AI agent capable of sensing and taking simple actions in the environment. The detail control and physical aspects of the system will be abstracted to a certain degree in the course. In future offerings experimental component with a real system is likely to be added but is beyond scope in the current offering.

References

- Mykel Kochenderfer, Decision Making Under Uncertainty
- Stuart Russell and Peter Norvig. Artificial Intelligence: A Modern Approach. 3rd Edition
- Sebastian Thrun, Wolfram Burgard and Dieter Fox. Probabilistic Robotics. MIT Press, 2005.
- Rich Sutton and Andrew Barto. Reinforcement Learning. MIT Press
- Steven LaValle. Planning Algorithms. Cambridge University Press, 2006.
- Dimitri P. Bertsekas. Dynamic Programming and Optimal Control. Athena Scientific, 3rd edition, 2007.
- Paper references will be added in due course.