ECE-GY9243 Optimal and Learning Control for Robotics

Syllabus

Description What kind of movements should a robot perform in order to walk, jump or manipulate objects? Can it compute optimal behaviors online? Can it learn this directly from trial and error? This course will introduce modern methods for robotics movement generation based on numerical optimal control and reinforcement learning. It will cover fundamental topics in numerical optimal control (Bellman equations, differential dynamic programming, model predictive control) and reinforcement learning (actor-critic algorithms, model-based reinforcement learning, deep reinforcement learning) applied to robotics. It will also contain hands-on exercises for real robotic applications such as walking and jumping, object manipulation or acrobatic drones.

Objective Students will learn modern methods for robotic motion planning and control based on numerical optimal control and reinforcement learning. They will also learn how to use these tools to solve practical robotics problems.

Topics Covered

1. Markov decision processes, Bellman equations
2. LQ problems with constraints, linear model predictive control in robotics
3. Iterative algorithms for non-linear problems (DDP and variants)
4. Trajectory optimization (movement planning as an optimization problem)
5. Fundamentals of reinforcement learning, Q-learning
6. Policy Improvement methods in RL (actor-critic, policy gradients)
7. Deep reinforcement learning (end-to-end RL, guided policy search, deep Q-learning)

Recommended background in at least one of the following: linear systems; robotics; machine learning; convex optimization; programming (python or C++)

Prerequisites Linear Systems (ECE-GY 6253 or ME-GY 6703) or Simulation tools for robotics ME-GY 6923