

CS592 Introduction to Robot Motion

Course Description

Autonomous robots will soon play a significant role in various domains, such as search-and-rescue, agriculture farms, homes, offices, transportation, and medical surgery, where fast, safe, and optimal response to different situations will be critical. However, to do so, these robots need fast algorithms to plan their motion sequences in real-time with limited perception and battery life. The field of motion planning and control addresses this challenge of coordinating robot motions and enabling them to interact with their environments for performing various challenging tasks under constraints. This class aims to introduce state-of-the-art methods in solving robot motion planning problems that emerged from the cross-fertilization of two long-lasting machine learning and motion planning fields.

The first part of this course focuses on the classical techniques in motion planning, i.e., RRT, RRT*, and their variants. The second part focuses on the modern approaches that bring together motion planning and machine learning. This course will cover basics in robot motion planning such as robot configuration space, collision avoidance, and it will evolve towards modern machine learning-based planning algorithms.

The syllabus is subject to change and modifications will be announced appropriately.

Time & Location:

TTH 3:00-4:15 pm, in LWSN B134

Prerequisites:

Students are expected to have background in Algorithm Analysis (Equivalent to CS381) and a strong programming background.

Administrative Information

Instructor: Ahmed Qureshi; qureshi7@purdue.edu

Teaching Assistant: TBD

Textbooks

There is no textbook for this course due to the emphasis on current techniques. Some helpful background references include:

- **Principles of Robot Motion** (Theory, Algorithms, and Implementations) by Howie Choset, Kevin M. Lynch, Seth Hutchinson, George A. Kantor, Wolfram Burgard, Lydia E. Kavraki and Sebastian Thrun.
- **Planning Algorithms** by LaValle, Steven M. Cambridge university press, 2006.
- **Deep learning** by Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. MIT press, 2016.

Grading

There are no midterm and final exams.

- Class participation & in-class discussion: 15%
- Paper reading and summary: 30%
- Paper presentation: 30%
- Final survey report: 25%

Paper Presentation: For each class, a student will lead the discussion by preparing:

- 1) a two-page summary per paper which will be electronically available to other students before class.
- 2) A 30-40 min presentation covering background on the topic and the main content of the assigned paper(s).

Depending on the number of students enrolled in the course, each student will present multiple times over the course of the semester.

Paper summary: It will be a 2-page, single column summary per paper.

Early Submission Deadline: Monday or Wednesday, i.e., a day before class, at 11:59:00 PM

Late Submission Deadline: Tuesday or Thursday, i.e., same day before class, at 10:00:00 AM with 10% penalty.

Survey report: Each student will submit a final survey report using IEEE Latex template. Deadline: December 16, 2021, 11:59 pm (EST)

Papers List

The paper list is preliminary and subject to change.

Lecture 1 (8/24): Intro to Robot Motion Planning (Lecture by the Instructor)

Lecture 2 (8/26): Sampling-based Motion Planners

- Rapidly exploring random trees: A new tool for path planning [[pdf](#)]
- RRT-connect: An efficient approach to single-query path planning [[pdf](#)]

Lecture 3 (8/31): Sampling-based Motion Planners

- Probabilistic Roadmaps for Path Planning in High Dimensional Configuration Spaces [[pdf](#)]

Lecture 4 (9/2): Optimal Sampling-based Motion Planning

- Incremental sampling-based algorithms for optimal motion planning [[pdf](#)]

Lecture 5 (9/7): Limitations of Sampling-based Methods, Advance Techniques

- Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic [\[pdf\]](#)

Lecture 6 (9/9): Limitations of Sampling-based Methods, Advance Techniques

- Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions [\[pdf\]](#)

Lecture 7 (9/14): Limitations of Sampling-based Methods, Advance Techniques

- Intelligent bidirectional rapidly-exploring random trees for optimal motion planning in complex cluttered environments [\[pdf\]](#)

Lecture 8 (9/16): Introduction to Deep Learning (Lecture by the Instructor)

Lecture 9 (9/21): Learning-based Motion Planning

- A robot path planning framework that learns from experience [\[pdf\]](#)

Lecture 10 (9/23): Learning-based Motion Planning

- LEGO: Leveraging Experience in Roadmap Generation for Sampling-Based Planning [\[pdf\]](#)

Lecture 11 (9/28): Learning-based Motion Planning

- Learning sampling distributions for robot motion planning [\[pdf\]](#)

Lecture 12 (9/30): Learning-based Motion Planning

- Data-driven planning via imitation learning [\[pdf\]](#)

Lecture 13 (10/5): Learning-based Motion Planning

- Motion planning networks: Bridging the gap between learning-based and classical motion planners [\[pdf\]](#)

Lecture 14 (10/7): Learning-based Motion Planning

- Learning Obstacle Representations for Neural Motion Planning [\[pdf\]](#)

10/12: Break

Lecture 15 (10/14): Learning-based Motion Planning

- Task-relevant Representation Learning for Networked Robotic Perception [\[pdf\]](#)

Lecture 16 (10/19): Learning-based Motion Planning

- Robot motion planning in learned latent spaces [\[pdf\]](#)

Lecture 17 (10/21): Learning-based Motion Planning

- Cost-to-Go Function Generating Networks for High Dimensional Motion Planning [\[pdf\]](#)

Lecture 18 (10/26): Learning-based Motion Planning

- Motion Planning Transformers: One Model to Plan Them All [\[pdf\]](#)

Lecture 19 (10/28): Learning-based Motion Planning

- Deep visual reasoning: Learning to predict action sequences for task and motion planning from an initial scene image [\[pdf\]](#)

Lecture 20 (11/2): Learning-based Motion Planning

- Harnessing reinforcement learning for neural motion planning [\[pdf\]](#)

Lecture 21 (11/4): Advance Classical Planning

- Eb-rrt: Optimal motion planning for mobile robots [\[pdf\]](#)

Lecture 20 (11/2): Learning-based Motion Planning

- RL-RRT: Kinodynamic motion planning via learning reachability estimators from RL policies [\[pdf\]](#)

Lecture 22 (11/9): Learning-based Motion Planning

- Long-range indoor navigation with prm-rl [\[pdf\]](#)

Lecture 23 (11/11): Learning-based Motion Planning:

- Path Planning for Manipulation Using Experience-Driven Random Tree [\[pdf\]](#)

Lecture 24 (11/16): Learning-based Motion Planning

- Learning latent dynamics for planning from pixels [\[pdf\]](#)

Lecture 25 (11/18): Learning-based Motion Planning

- MPC-MPNet: Model-Predictive Motion Planning Networks for Fast, Near-Optimal Planning under Kinodynamic Constraints [\[pdf\]](#)

Lecture 26 (11/23): Learning-based Motion Planning

- Path planning with local motion estimations [\[pdf\]](#)

Lecture 27 (11/30): Learning-based Motion Planning

- Learning When to Trust a Dynamics Model for Planning in Reduced State Spaces [\[pdf\]](#)

Lecture 28 (12/2): Learning-based Motion Planning

- Sampling-based methods for motion planning with constraints [\[pdf\]](#)

Lecture 29 (12/7): Learning-based Motion Planning

- Constrained motion planning networks X [\[pdf\]](#)

Lecture 30 (12/9): Learning-based Motion Planning

- The Blindfolded Robot: A Bayesian Approach to Planning with Contact Feedback [\[pdf\]](#)