MO89905: Learning-Based Control for Mobility Systems

Fall 2026

L:L:C	3:0:3	**Course Type**	Elective (Graduate)
Class Time	T/Th 13:00–14:30	**Class Room**	L409, Munji Campus
Instructor	Kyunghwan Choi	**Email**	kh.choi@kaist.ac.kr
Office Hours	After class or by appointment	**Course Page**	Online textbook

Description: This course develops learning-based control through exact and approximate dynamic programming (DP). Approximate DP provides a common framework for interpreting reinforcement learning and connecting learned values, policies, and models to LQR, MPC, and rollout. A central emphasis is online improvement: using lookahead and optimization to improve a learned controller during execution. Modern RL algorithms are then interpreted within this framework rather than treated as isolated recipes.

Learning Objectives: By the end of the course, students should be able to:

- formulate finite- and infinite-horizon DP problems and explain how value, policy, and model approximation lead to RL methods;
- connect approximate DP with LQR, MPC, and rollout, and design online improvement using learned values, policies, or models with lookahead;
- design and evaluate mobility controllers using credible baselines, shifted conditions, control metrics, timing, and reproducible evidence, while interpreting modern RL within the same framework.

Textbook: K. Choi, *Learning-Based Control for Mobility Systems*, evolving working draft, 2026.
https://kyunghwan-choi.github.io/learning-based-control-mobility/

References:

- D. Bertsekas, *Reinforcement Learning and Optimal Control*, 2019.
- D. Bertsekas, *A Course in Reinforcement Learning*, 2nd ed., 2024.
- R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, 2nd ed., 2018.

Prerequisites: Linear algebra, calculus, probability, basic feedback control, and MATLAB or Python programming. Prior RL coursework is not required.

Assessment:

Homework	20%	Slalom design challenge	35%
Term project	35%	Attendance	10%

Slalom design challenge. This replaces a conventional written midterm. Students formulate, implement, and evaluate a learning-based controller for a provided slalom problem.

Term project. Students address an approved learning-based mobility control problem and submit both a technical report and an oral presentation.

Course Policy: The course follows KAIST rules. Sources, reused code, collaboration, and material use of AI tools must be disclosed, and submitted experiments must be reproducible.

Tentative Course Schedule

Week	Topic	Remark
1	Introduction: optimal control, DP, RL, and online improvement	HW1 released
2	Dynamic programming foundations: finite and infinite horizons	—
3	Exact and approximate DP	HW1 due; HW2 released
4	Parametric approximation in value, policy, and model space	Slalom challenge released
5	Approximation in value space and policy evaluation	HW2 due
6	Rollout and online lookahead	—
7	LQR and MPC through the DP lens	Slalom checkpoint
8	Slalom design review	Slalom challenge due
9	Approximation in policy space and actor–critic	HW3 released; project topics
10	Modern RL through value and policy approximation	—
11	Online policy improvement	HW3 due; HW4 released
12	Model learning	Project proposal due
13	Control with learned models	—
14	Safety and online adaptation	HW4 due
15	Multi-agent coordination and course synthesis	Project report due
16	Project presentations	Presentation

The schedule may be adjusted to the Fall 2026 academic calendar. Timetable, room, course page, and assignment dates will be confirmed before release. Updated July 16, 2026.