COURSE: ECE 780T - Topic 13 (Fall 2026)

COURSE TITLE: **Robot Learning**

INSTRUCTOR: Roya Firoozi roya.firoozi@uwaterloo.ca

DESCRIPTION:

This course introduces robot learning, focusing on data-driven and foundation-model approaches for enabling generalizable robotic autonomy. The course begins with the foundations of robot learning and generative modeling and progresses to imitation learning and vision-language-action (VLA) models for large-scale robotic systems. Core learning paradigms, including reinforcement learning, are introduced alongside model-based control approaches to provide foundational perspectives on decision-making and control in robotics. Additional topics include adaptation and fine-tuning of robotic models, representation learning for perception and planning, multi-agent learning, and uncertainty quantification and safety considerations for AI-driven robotic systems. Overall, the course emphasizes scalable, adaptable, and reliable approaches for developing intelligent robotic systems capable of operating in complex and unstructured environments.

COURSE/TEACHING OBJECTIVES:

This course will enable students to:

- Develop a rigorous understanding of the theoretical foundations of robot learning, including generative modeling, imitation learning, reinforcement learning and optimal control, representation learning, multi-agent learning, adaptation, safety analysis.
- Analyze and compare learning algorithms in terms of assumptions, convergence properties, sample efficiency, generalization, and computational complexity.
- Formulate robotic decision-making and control problems within principled mathematical frameworks (e.g., generative trajectory modeling, MDPs, probabilistic graphical models, and optimal control).
- Design and implement modern learning algorithms for robotic systems and evaluate their empirical performance.
- Critically read, interpret, and synthesize contemporary research literature in robot learning.
- Identify open research challenges and articulate informed perspectives on current trends and limitations in the field.
- Communicate technical ideas effectively through written reports and oral presentations.

SYLLABUS:

01. Introduction to Robot Learning
    Motivation and examples, classical control vs. learning-based robotics, learning paradigms, challenges and open problems, course roadmap
02. Introduction to Generative Modeling
    Definition and probabilistic foundations, sequential generative models, overview of generative model families (autoregressive models, variational autoencoders, diffusion models, and flow-based models), connections to robotics
03. Imitation Learning
    Behavioral cloning, flow matching, diffusion policy, conditioning and guidance
04. Robot Transformers and Vision-Language-Action (VLA) Models
    Tokenization of robotic data, transformer policies, architectures and examples of VLA models
05. Value-Based and Policy Gradient Methods
    Markov decision processes (MDPs), Q-learning, policy gradient theorem, actor–critic methods, advanced policy optimization
06. Model-Based Control
    Model Predictive Control (MPC)
07. Adaptation and Fine-Tuning
    Transfer learning, fine-tuning
08. World Models and Spatial Representations for Robotics
    World models, planning with 3D visual-language fields
09. Special Topic: Multi-Agent Learning
    Multi-agent behavioral cloning, multi-agent reinforcement learning
10. Special Topic: Uncertainty Quantification and Safety
    Black-box statistical methods for uncertainty quantification in large AI-driven robotics models

TEXTBOOK:

No textbook required. Lecture slides will be provided. The slides are intended as a supplement to the lectures and may not include all material covered in class. Additional explanations, derivations, and examples will be presented on the board during lectures, and students are expected to take notes accordingly. Throughout the course, relevant peer-reviewed papers will be posted on the course webpage.

RECOMMENDED BACKGROUND:

There are no formal prerequisites for this course. Some background in machine learning, optimization, probability, and control theory is recommended.

GENERAL REFERNCES:

- "Reinforcement Learning: An Introduction", Richard S.Sutton and Andrew G. Barto, MIT Press, 2nd edition (2018).
- "Deep Learning", Ian Goodfellow and Yoshua Bengio and Aaron Courville, MIT Press, (2016).

MARKING SCHEME:

- **30% Homework**
  There will be three homework assignments throughout the term.
- **20% Midterm Exam**
  The exam will be approximately 80 minutes long, will be held during the regular lecture time, and is scheduled for the midterm week.
- **35% Final Project**
  Students will complete a project on a topic related to robot learning. The project may be completed individually or in groups of two.
- **10% Research Paper Presentation**
  - Students may complete the presentation individually or in groups of two.
  - Each presentation will be 15 minutes, followed by 5 minutes of questions and discussion.
  - The presentation will be based on a peer-reviewed research paper selected from a list of papers provided on the LEARN course webpage.
  - A sign-up sheet will be made available, and presentation dates will be assigned on a first-come, first-served basis. Presentations will be scheduled throughout the term.
  - Research paper presentations are separate from and independent of the final project presentations.
- **5% Participation**
  Participation in weekly question-and-answer sessions and class discussions.

COURSE WEBSITE:

A course homepage is available on LEARN.

TENTATIVE SCHEDULE:

A tentative schedule is shown below. The instructor reserves the right to change the outline and/or the schedule as needed.

Week No.	Main Subject/Topics Covered During Lecture
1	1. Intro to robot learning
2	2. Intro to generative modeling I 3. Intro to generative modeling II
3	4. Imitation learning I 5. Imitation learning II
4	6. Robot transformers, VLA I 7. Robot transformers, VLA II
5	8. MDPs, value-based methods 9. Policy gradient methods
6	10. Actor-critic methods 11. Advanced policy optimization
7	Midterm
8	12. Model-based control: MPC 13. Adaptation and fine-tuning
9	14. World models and spatial representations for robotics 15. Special topic: multi-agent learning I
10	16. Special topic: multi-agent learning II 17. Special topic: uncertainty quantification and safety I
11	18. Special topic: uncertainty quantification and safety II 19. Final project oral presentations
12	20. Final project oral presentations 21. Final project oral presentations

FINAL PROJECT:

The goal of the final project is to explore a topic in robot learning that builds on concepts covered in the course (e.g., reinforcement learning, diffusion policies, imitation learning, planning, generative models, world models, or multi-agent systems, robot safety, VLAs). Students will design, implement, and evaluate a method or system that demonstrates learning-based decision-making for robotics.

Group Formation

Students may complete the final project individually or in groups of two. Students who choose to work in groups are expected to form their own groups.

Final Project Scope

Projects should:

- Focus on robot learning (broadly defined)
- Be connected to course topics
- Include experimental evaluation (simulation or real-world)

Example directions (not exhaustive):

- Reinforcement learning for control or planning
- Diffusion / generative policies for robot actions
- Vision-language models for robot decision-making
- Learning-based navigation or manipulation
- Multi-agent or human-robot interaction
- Learning with uncertainty or safety constraints
- Combining classical planning/control with learning
- Fine tunning

Expectations

A strong project should include:

- A clear and novel problem formulation (introduce a clear contribution beyond existing literature)
- A well-motivated approach
- Implementation (from scratch or built on existing frameworks)
- Quantitative evaluation (baselines are strongly encouraged)
- Clear analysis of results and limitations

Deliverables

1. Written Report

- Format: IEEE conference style (double column)
- Length: 6 pages
- Should include:
- Abstract
- Introduction & related work
- Methodology
- Experimental setup
- Results & discussion
- Conclusion

2. Oral Presentation

- Duration: 20 minutes + Q&A

- Should clearly cover:
- Problem motivation
- Key idea
- Method overview
- Results and insights
- Limitations and future work

Milestones

1. Proposal (1–2 pages)
   Problem, approach, expected contribution

2. Midterm check-in
   Preliminary results or progress update

3. Final submission
   Written report + oral presentation

Evaluation Criteria

Projects will be evaluated based on:

- Technical depth (30%)
- Clarity and presentation (20%)
- Experimental validation (25%)
- Originality / insight (15%)
- Relevance to robot learning (10%)