ROBOT LEARNING SEMINAR

CMSC848X · Fall 2026 · Computer Science · University of Maryland, College Park

Instructor	Ruoshi Liu	Meeting time	Tue & Thu, [14:00–15:15]
Email	ruoshi@umd.edu (currently unavailable, email ruoshi98@gmail.com instead)	Location	CSI 1122
Office hours	IRB 4218, Thu 15:30-16:30	Credits	3
Teaching assistant	George Zhang qz2002@umd.edu		

Course slack:
https://join.slack.com/t/robotlearning-vty9190/shared_invite/zt-485209wkl-KU9PKmLsEfSQ5SvVrHpZEA

1. Course Description

This course introduces the core methods that let robots learn from data. We focus on visuomotor policy learning for manipulation, covering behavior cloning, offline and online reinforcement learning, sim-to-real transfer, and multimodal perception across vision, touch, sound, and force. Alongside the methods, we take up the problems that make robotics distinctive: embodiment, data efficiency, and generalization across tasks and platforms — and the directions currently reshaping the field, including robotic foundation models and world models.

The course is run as a seminar built on recent research papers. There are no textbook lectures after the first week. Instead, students present, argue, and defend papers in structured roles, then discuss them in small groups. The semester ends with a research-style final project.

Learning objectives

- Read a robotics or robot-learning paper critically and quickly, separating what was demonstrated from what was claimed.
- Explain the dominant families of visuomotor policy learning — behavior cloning, RL, world models, foundation models — and articulate when each is the right tool.
- Argue both sides of a paper: build the strongest case for it and against it.
- Identify the specific engineering and algorithmic decisions that make robot learning systems actually work in the real world.

- Design, execute, and report a small original research project in robot learning.

Prerequisites

Graduate standing or instructor permission. Students should have prior coursework or equivalent experience in machine learning and deep learning, comfort with linear algebra and probability, and working fluency in Python and PyTorch. Prior robotics coursework is helpful but not required; prior experience with a physical robot is not assumed.

2. How the Seminar Works

The week has a fixed rhythm. Tuesdays are Topic Days and Thursdays are Researcher Days. Every session runs on the same 75-minute clock:

TIME	SEGMENT	WHAT HAPPENS
45 min	Presentation	Student teams present, in role. Topic Days use the eight-role structure below; Researcher Days are group paper presentations.
15 min	Breakout discussion	The class splits into groups of ~7 to discuss the papers away from the podium.
15 min	Conclusion	Groups report back; the presenting team synthesizes and closes.

Tuesday — Topic Day

Each Topic Day covers one theme through two or three papers. The presenting team opens with two framing slides — one on the topic as a whole, one situating the specific papers — and then splits into eight roles. Each role owns roughly one slide and five minutes.

The authors and reviewers are adversarial by design. The author's job is to sell the paper as hard as they honestly can. The reviewer's job is to reject it. Neither role is a summary; both are arguments. The learners, on the other hand, are neutral summarizers who try to discover high-level trends and ideas as well as low-level but crucial technical decisions.

ROLE	MANDATE
Task author	Frame the problem the paper attacks. Why does it matter, why is it hard, and what does solving it buy the field? Sell it.
Task reviewer	Argue the problem is ill-posed, already solved, narrower than claimed, or not worth the field's attention.
Approach author	Present the method and defend every design choice as the right one under the constraints.
Approach reviewer	Attack the method: unjustified choices, hidden assumptions, brittle components, complexity without payoff.
Result author	Present the evidence and argue it supports the paper's claims. Highlight the experiments that actually matter.
Result reviewer	Attack the evaluation: weak baselines, cherry-picked tasks, missing ablations, generalization claims the data cannot bear.

ROLE	MANDATE
High-level learner	What are the high-level ideas? How do they differ from prior work? What does this open up for future work?
Low-level learner	What is the key trick that made it work? What smart ML / robotics / algorithmic techniques were used? What would you do differently?

The deck also carries three closing slides: one synthesizing the week's Slack thread, one posing the breakout prompts, and one conclusion drawn from what the breakout groups reported. Total deck: 2 framing + 8 role + 3 closing slides.

Thursday — Researcher Day

Thursdays are built around a single researcher rather than a single topic. Each group picks one paper by that researcher and presents it. Three rules:

- The researcher must be either first or last author on the paper you choose.
- No survey papers.
- First come, first served — claim your paper in Slack by **Monday 10:00 pm**. No two groups may present the same paper.

Group presentations follow the Task / Approach / Result spine from Topic Days, without the adversarial split, and close with one high-level takeaway and one low-level trick. Each group has 5 minutes. The point of the day is to see how one researcher's questions, methods, and taste evolve across a body of work — so tell us how your paper fits that arc.

Slack discussion

Before every class, each student posts one substantive question about the assigned papers or researchers and responds to one question posted by someone else. Both are due by **10:00 pm the day before class**. A good question is one you could not answer by just reading the abstract. The presenters can take inspirations from the slack discussions.

Breakout discussion

The class splits into groups of roughly seven. Groups are re-shuffled periodically so you are not always arguing with the same people. Discuss any of:

- Something new you learned from the paper.
- An insight you formed while reading it.
- An idea or research direction it inspired.
- A reflection worth sharing with the rest of the class.

Each group designates a reporter who has two minutes in the conclusion segment. The presenting team closes the session by synthesizing what came out of the room.

3. Final Project

The final project is a small piece of original research in robot learning, carried out in teams of 2-4. It should look like a workshop paper: a clear question, an honest experiment, and a result you can defend — including a negative one. Reproducing a paper is acceptable only if you extend it with a question the original did not answer.

Projects may be simulation-only. Access to physical hardware is not required. Compute is available through university cluster.

MILESTONE	DUE	WHAT TO SUBMIT
Team formation + topic sketch	Thur, Sep 17	One paragraph and a team roster, posted in Slack. Ungraded but required.
Project proposal (1–2 pages)	Thur, Oct 8	Question, related work, proposed method, evaluation plan, risks and fallback.
Mid-project check-in (1 page)	Thur, Nov 5	What works, what broke, what changed. Paired with a required 15-minute meeting.
Final presentation	Dec 8 & 10	[12–15] min per team plus questions, in class.
Final report + code	Thur, Dec 10	4–8 pages in a standard conference format, plus a runnable repository.

4. Grading

COMPONENT	WEIGHT	NOTES
Slack discussion	10%	One question and one reply before each class. Graded on engagement and substance, not volume.
Course presentations	50%	Topic Day role performance (30%) + Researcher Day group presentation (20%).
Final project	40%	Proposal 5% · check-in 5% · final presentation 15% · report and code 15%.

Final project rubric

CRITERION	WEIGHT	WHAT EXCELLENT LOOKS LIKE
Question and motivation	20%	The question is sharp, unresolved by prior work, and worth a semester.
Technical execution	30%	Sound method, correct implementation, sensible baselines, controlled comparisons.
Evidence and analysis	25%	Claims are supported by the experiments run. Ablations isolate the thing you say matters. Negative results reported honestly.
Communication	20%	Report and talk are clear, well-figured, and appropriately scoped.
Reproducibility	5%	Code runs from a clean checkout with documented setup.

Grade scale: A 90-100, B 80-90, C 70-80.

5. Course Policies

Use of large language models

The intellectual work of this seminar is forming and defending your own judgment about research papers. An LLM cannot do that for you. However, we do recognize the benefits and the necessity of using LLMs in today's research and studies, therefore, the following rules will apply:

Permitted:

- Polishing language in text you wrote yourself — grammar, phrasing, translation.
- Explaining background concepts, notation, or mathematics you encounter while reading.
- Debugging your own code, or visualization.

Not permitted:

- Generating slide content, paper summaries, critiques, discussion questions, or Slack posts.
- Using an LLM to summarize a paper in place of reading it.
- Generating project ideas, research questions, experimental designs.
- Submitting model-generated text or code as your own work.

The rule in one line: the ideas have to be yours; polishing the language is fine. Disclose any LLM use in a short note at the end of your report or code — naming the tool and what you used it for. Disclosure is never penalized; undisclosed use is an academic integrity violation.

Attendance and participation

This is a discussion course; attendance is expected and participation is graded through both the in-person course and slack discussions. Notify the instructor in advance of an unavoidable absence. If you are presenting and cannot attend, swap with your friend and tell the instructor at least 48 hours ahead. You lose 1 point for missing a course, and 2 points for missing a presentation.

Technology and accessibility

Presenting teams need a laptop for slides, and all students need reliable access to Slack. All readings in this course are freely available; there is no textbook to buy.

If a laptop or reliable internet is a barrier for you at any point in the semester, tell the instructor. We can figure out some potential ways to gain access such as equipment loaning program from library or department.

Accommodations

Students with documented disabilities should contact Accessibility and Disability Service as early in the semester as possible. Accommodations will be implemented in full, including for presentation format and timing. You do not need to disclose the nature of a disability to the instructor.

Academic integrity

All work must be your own or your team's, with contributions from others — collaborators, papers, code repositories, and AI tools — clearly credited. Reusing figures or text from a paper without attribution, or presenting another group's analysis as yours, is a violation.

Late work

Slack posts are not accepted late — the thread has to exist before class to be useful. Extensions for illness or emergency are granted freely with official notice.

6. Semester Schedule

Tuesdays are Topic Days; Thursdays are Researcher Days. Weeks 1 and 15 and the two sessions stranded by Fall Break and Thanksgiving are instructor lectures, guest lectures, and final presentations.

WEEK	TUESDAY — TOPIC DAY	THURSDAY — RESEARCHER DAY
1 Aug 31–Sep 4	TUE, SEP 1 · INSTRUCTOR LECTURE Course Overview & Logistics Instructor lecture. What the seminar is, how the roles work, how to read a robotics paper, sign-ups.	THU, SEP 3 · INSTRUCTOR LECTURE Foundations of Robot Learning Instructor lecture. Problem setup, embodiment, data, evaluation — the vocabulary for the rest of the term.
2 Sep 7–11 *Labor Day Mon, Sep 7*	TUE, SEP 8 · TOPIC DAY Visual Representation for Robotics DINOv1 · DINOv2 · DINOv3	THU, SEP 10 · RESEARCHER DAY Jitendra Malik Group-selected papers
3 Sep 14–18	TUE, SEP 15 · TOPIC DAY Behavior Cloning IBC · ACT	THU, SEP 17 · RESEARCHER DAY Shuran Song Group-selected papers
4 Sep 21–25	TUE, SEP 22 · TOPIC DAY Reinforcement Learning DPPO · TD-MPC2	THU, SEP 24 · RESEARCHER DAY Richard Sutton Group-selected papers
5 Sep 28–Oct 2	TUE, SEP 29 · TOPIC DAY Simulation & Sim-to-Real Transfer SimToolReal · RMA	THU, OCT 1 · RESEARCHER DAY Xue Bin Peng Group-selected papers
6 Oct 5–9	TUE, OCT 6 · TOPIC DAY World Models Curiosity · Interactive World Simulator	THU, OCT 8 · RESEARCHER DAY Danijar Hafner Group-selected papers
7 Oct 12–16 *Fall Break Mon–Tue*	TUE, OCT 13 *Fall Break — no class*	THU, OCT 15 · GUEST LECTURE Guest Lecture Huy Ha (Anthropic)
8 Oct 19–23	TUE, OCT 20 · TOPIC DAY Language Models in Robotics π0.7 · VoxPoser	THU, OCT 22 · RESEARCHER DAY Andy Zeng Group-selected papers
9 Oct 26–30	TUE, OCT 27 · TOPIC DAY Multisensory & Multimodal Perception See, Hear, and Feel · ManiWAV	THU, OCT 29 · RESEARCHER DAY Andrew Owens Group-selected papers
10 Nov 2–6	TUE, NOV 3 · TOPIC DAY Tactile Manipulation Reactive Diffusion Policy · 3D-ViTac	THU, NOV 5 · RESEARCHER DAY Roberto Calandra Group-selected papers

WEEK	TUESDAY — TOPIC DAY	THURSDAY — RESEARCHER DAY
11 Nov 9–13	TUE, NOV 10 · TOPIC DAY Grasping Dexonomy · SPIDER	THU, NOV 12 · RESEARCHER DAY Ken Goldberg Group-selected papers
12 Nov 16–20	TUE, NOV 17 · TOPIC DAY Dexterous Manipulation DextrAH-RGB · Visual Dexterity	THU, NOV 19 · RESEARCHER DAY Deepak Pathak Group-selected papers
13 Nov 23–27 *Recess begins Wed, Nov 25*	TUE, NOV 24 · GUEST LECTURE Guest Lecture Daniel Seita (USC)	THU, NOV 26 *Thanksgiving — no class*
14 Nov 30–Dec 4	TUE, DEC 1 · TOPIC DAY Computational Hardware Design DGDM · Hardware as Policy	THU, DEC 3 · RESEARCHER DAY Wojciech Matusik Group-selected papers
15 Dec 7–11 *Classes end Fri, Dec 11*	TUE, DEC 8 · FINAL PRESENTATIONS Final Project Presentations I	THU, DEC 10 · FINAL PRESENTATIONS Final Project Presentations II
Finals Dec 14 – 21	No final exam for this course. Good luck with other exams!	

7. Topics, Researchers, and Readings

Every paper below is free on arXiv. Read the assigned papers before the Tuesday session; Thursday papers are chosen by your group.

TOPIC & RESEARCHER	PAPERS
Visual Representation for Robotics **Thursday: Jitendra Malik**	· Emerging Properties in Self-Supervised Vision Transformers (DINO) arXiv:2104.14294 · DINOv2: Learning Robust Visual Features without Supervision arXiv:2304.07193 · DINOv3 arXiv:2508.10104
Behavior Cloning **Thursday: Shuran Song**	· Implicit Behavioral Cloning (IBC) arXiv:2109.00137 · Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware (ACT) arXiv:2304.13705
Reinforcement Learning **Thursday: Richard Sutton**	· Diffusion Policy Policy Optimization (DPPO) arXiv:2409.00588 · TD-MPC2: Scalable, Robust World Models for Continuous Control arXiv:2310.16828
Simulation & Sim-to-Real Transfer **Thursday: Xue Bin Peng**	· SimToolReal: An Object-Centric Policy for Zero-Shot Dexterous Tool Manipulation arXiv:2602.16863 · RMA: Rapid Motor Adaptation for Legged Robots arXiv:2107.04034
World Models **Thursday: Danijar Hafner**	· Curiosity-driven Exploration by Self-supervised Prediction arXiv:1705.05363 · Interactive World Simulator for Robot Policy Training and Evaluation arXiv:2603.08546
Language Models in Robotics **Thursday: Andy Zeng**	· π0.7: a Steerable Generalist Robotic Foundation Model with Emergent Capabilities arXiv:2604.15483 · VoxPoser: Composable 3D Value Maps for Robotic Manipulation with Language Models arXiv:2307.05973
Multisensory & Multimodal Perception **Thursday: Andrew Owens**	· See, Hear, and Feel: Smart Sensory Fusion for Robotic Manipulation arXiv:2212.03858 · ManiWAV: Learning Robot Manipulation from In-the-Wild Audio-Visual Data arXiv:2406.19464
Tactile Manipulation **Thursday: Roberto Calandra**	· Reactive Diffusion Policy: Slow-Fast Visual-Tactile Policy Learning for Contact-Rich Manipulation arXiv:2503.02881 · 3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing arXiv:2410.24091
Grasping **Thursday: Ken Goldberg**	· Dexonomy: Synthesizing All Dexterous Grasp Types in a Grasp Taxonomy arXiv:2504.18829 · SPIDER: Scalable Physics-Informed Dexterous Retargeting arXiv:2511.09484
Dexterous Manipulation **Thursday: Deepak Pathak**	· DextrAH-RGB: Visuomotor Policies to Grasp Anything with Dexterous Hands arXiv:2412.01791 · Visual Dexterity: In-Hand Reorientation of Novel and Complex Object Shapes arXiv:2211.11744
Computational Hardware Design **Thursday: Wojciech Matusik**	· Dynamics-Guided Diffusion Model for Sensor-less Robot Manipulator Design arXiv:2402.15038 · Hardware as Policy: Mechanical and Computational Co-Optimization using Deep RL arXiv:2008.04460