

CS577: Robot Learning and Interaction

로봇 학습과 상호 작용

Instructor

Asst. Prof. Daehyung Park (daehyung@kaist.ac.kr)

Office hours: Anytime by appointment

Time & Location

Tue/Thur 10:30 - 12:00 at E3-1 2443

Teaching Assistant

Dohyun Kim (dohyun141@kaist.ac.kr)

Description

This course will introduce graduate students to the emerging area of intelligent robotics toward human-centered robotics. The course overviews each robotic learning and interaction area including learning from demonstration (LfD), reinforcement learning (RL), natural language interaction, interactive perception, etc. We will then review the state-of-the-art technologies, particularly outstanding award/finalist papers, and exercise a part of technologies using simulated robotic manipulators. Finally, we will exercise the learned techniques via final individual/team projects.

Prerequisites:

It is recommended to take CS376, CS470, or CS477 before or with this class.

Schedule

# of lecture, date	Topics and slides
Sept - 2 (Tue.)	Survey for Robot Learning [1] Kroemer, Oliver, Scott Niekum, and George Konidaris. " A review of robot learning for manipulation: Challenges, representations, and algorithms ," arXiv preprint arXiv:1907.03146 (2019). [2] Roy, Nicholas, Ingmar Posner, Tim Barfoot, Philippe Beaudoin, Yoshua Bengio, Jeannette Bohg, Oliver Brock et al. " From Machine Learning to Robotics: Challenges and Opportunities for Embodied Intelligence ," arXiv preprint arXiv:2110.15245 (2021).
Sept - 4 (Thur.)	Learning from demonstrations I - No real-time lecture. Please, check the recorded video on KLMS [1] Brenna D. Argall, Sonia Chernova, Manuela Veloso, Brett Browning, " A survey of robot learning from demonstration ," <i>Robotics and Autonomous Systems (RAS)</i> , 2008. [2] Calinon, Sylvain. " Learning from demonstration (programming by demonstration) ," <i>Encyclopedia of Robotics</i> , 2018.
Sept - 9 (Tue.)	Learning from demonstrations II - Dynamic Movement Primitives [1] Pastor, Peter, Heiko Hoffmann, Tamim Asfour, and Stefan Schaal. " Learning and generalization of motor skills by learning from demonstration ," In <i>2009 IEEE International Conference on Robotics and Automation</i> , pp. 763-768. IEEE, 2009. [2] Hoffmann, Heiko, Peter Pastor, Dae-Hyung Park, and Stefan Schaal. " Biologically-inspired dynamical systems for movement generation: automatic real-time goal adaptation and obstacle avoidance ," In <i>2009 IEEE International Conference on Robotics and Automation</i> , pp. 2587-2592. IEEE, 2009.

Sept - 11 (Thur.)	Reserved for Guest Lecture: Dr. Gregory Joseph Stein from George Mason University
Sept - 16 (Tue.)	Learning from demonstrations III - Diffusion Policy [1] Chi, Cheng, et al. " Diffusion policy: Visuomotor policy learning via action diffusion ." <i>The International Journal of Robotics Research</i> (2023): 02783649241273668.
Sept - 18 (Thur.)	Learning from demonstrations IV - Vision Language Action (VLA) Models [1] Zhong, Yifan, et al. " A Survey on Vision-Language-Action Models: An Action Tokenization Perspective ." <i>arXiv preprint arXiv:2507.01925</i> (2025). Assignment I by Sept. 30th
Sept - 23 (Tue.)	Reinforcement Learning I - MDP [1] Kober, Jens, J. Andrew Bagnell, and Jan Peters. " Reinforcement learning in robotics: A survey ." <i>The International Journal of Robotics Research</i> 32.11 (2013): 1238-1274. (~ Sec. 2.1) [2] Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction . MIT press, 2018. (Sec. 3)
Sept - 25 (Thur.)	Reinforcement Learning II - Dynamic Programming [1] Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction . MIT press, 2018. (Sec. 3-4)
Sept - 30 (Tue.)	Reinforcement Learning III - Temporal Difference Learning [1] Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction . MIT press, 2018. (Sec. 5-6) [2] Mnih, V., Kavukcuoglu, K., Silver, D. <i>et al.</i> Human-level control through deep reinforcement learning . <i>Nature</i> 518 , 529–533 (2015).
Oct - 2 (Thur.)	Reinforcement Learning IV - Function Approximation [1] Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction . MIT press, 2018. (Sec. 13) Submit the title of your potential project by Oct 4th!
Oct - 7 (Tue.)	No class (Holiday)
Oct - 9 (Thur.)	No class (Holiday)
Oct - 14 (Tue.)	Reserved for Guest Lecture: Dr. Shen Li from CSAIL, MIT Assignment 2 by Oct. 28th
Oct - 16 (Thur.)	Inverse Reinforcement Learning I [1] Ziebart, Brian D., Andrew L. Maas, J. Andrew Bagnell, and Anind K. Dey. " Maximum entropy inverse reinforcement learning ." In <i>AAAI</i> , vol. 8, pp. 1433-1438. 2008. [2] Finn, Chelsea, Sergey Levine, and Pieter Abbeel. " Guided cost learning: Deep inverse optimal control via policy optimization ." International conference on machine learning. 2016.
Oct - 21 (Tue.)	No class (Midterm)
Oct - 23 (Thur.)	No class (Midterm)
Oct - 28 (Tue.)	Midterm project proposal (5-10min presentation + Q/A), problem+potential solution (+experiment plan) 1. 2. 3.
Oct - 30 (Thur.)	Mid term project proposal 1. 2. 3.

Nov - 4 (Tue.)	<p>[Student Presentation] Reinforcement Learning</p> <p>[1] Ankile, Lars, et al. "From imitation to refinement-residual rl for precise assembly." <i>2025 IEEE International Conference on Robotics and Automation (ICRA)</i>. IEEE, 2025.</p> <p>[2] Chen, Haonan, et al. "Tool-as-interface: Learning robot policies from human tool usage through imitation learning." <i>arXiv preprint arXiv:2504.04612</i> (2025).</p>
Nov - 6 (Thur.)	<p>[Student Presentation] Vision-language Action Models</p> <p>[1] Xu, Charles, et al. "Rldg: Robotic generalist policy distillation via reinforcement learning." <i>arXiv preprint arXiv:2412.09858</i> (2024).</p> <p>[2] Kang, Gi-Cheon, et al. "CLIP-RT: Learning Language-Conditioned Robotic Policies from Natural Language Supervision." <i>Robotics: Science and Systems</i>, 2025.</p>
Nov - 11 (Tue.)	<p>[Student Presentation] Vision-language Action Models</p> <p>[1] Kim, Moo Jin, Chelsea Finn, and Percy Liang. "Fine-tuning vision-language-action models: Optimizing speed and success." <i>Robotics: Science and Systems</i>, 2025.</p> <p>[2] Chen, Yuhui, et al. "Conrft: A reinforced fine-tuning method for vla models via consistency policy." <i>Robotics: Science and Systems</i>, 2025.</p>
Nov - 13 (Thur.)	<p>[Student Presentation] Humanoid Robots + ETC</p> <p>[1] Cheng, An-Chieh, et al. "Navila: Legged robot vision-language-action model for navigation." <i>Robotics: Science and Systems</i>, 2025.</p> <p>[2] Yu, Justin, et al. "Real2render2real: Scaling robot data without dynamics simulation or robot hardware." <i>arXiv preprint arXiv:2505.09601</i> (2025).</p>
Nov - 18 (Tue.)	<p>[Student Presentation] Multimodal Interaction</p> <p>[1] Huang, Binghao, et al. "VT-Refine: Learning Bimanual Assembly with Visuo-Tactile Feedback via Simulation Fine-Tuning." <i>CoRL 2025</i>.</p> <p>[2] Xiong, Haoyu, et al. "Vision in Action: Learning Active Perception from Human Demonstrations." <i>arXiv preprint arXiv:2506.15666</i> (2025).</p>
Nov - 20 (Thur.)	<p>[Student Presentation] Safety Learning</p> <p>[1] Zhang, Songyuan, et al. "Discrete GCBF proximal policy optimization for multi-agent safe optimal control." <i>International Conference on Learning Representations (ICLR)</i>, 2025.</p> <p>[2] Ni, Minheng, et al. "Don't Let Your Robot be Harmful: Responsible Robotic Manipulation via Safety-as-Policy." <i>arXiv preprint arXiv:2411.18289</i> (2024).</p>
Nov - 25 (Tue.)	<p>[Student Presentation] Safety Learning</p> <p>[1] Tölle, Maximilian, et al. "Towards Safe Robot Foundation Models Using Inductive Biases." <i>arXiv preprint arXiv:2505.10219</i> (2025).</p> <p>[2] Zhang, Borong, et al. "SafeVLA: Towards Safety Alignment of Vision-Language-Action Model via Constrained Learning." <i>arXiv preprint arXiv:2503.03480</i> (2025).</p>
Nov - 27 (Thur.)	No class (Undergrad interview)
Dec - 2 (Tue.)	<p>[Student Presentation] Task and/or motion planning</p> <p>[1] Wu, Yi, et al. "SELP: Generating safe and efficient task plans for robot agents with large language models." <i>arXiv preprint arXiv:2409.19471</i> (2024).</p> <p>[2] Pan, Tianyang, Rahul Shome, and Lydia E. Kavraki. "Task and motion planning for execution in the real." <i>IEEE Transactions on Robotics</i> 40 (2024): 3356-3371.</p>
Dec - 4 (Thur.)	<p>[Student Presentation] Human-Robot Interaction</p> <p>[1] Luo, Shengcheng, et al. "Human-agent joint learning for efficient robot manipulation skill acquisition." <i>arXiv preprint arXiv:2407.00299</i> (2024).</p> <p>[2] Lee, Hee Rin, et al. "Minding the Stop-Gap: Attending to the “Temporary.” Unplanned, and Added Labor of Human-Robot Collaboration in Context." <i>2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)</i>. IEEE, 2025.</p>
Dec - 9 (Tue.)	Final paper + presentation
Dec - 11 (Thur.)	Final paper + presentation
Dec - 16 (Tue.)	No class due to Final Exam
Dec - 18 (Thur.)	<p>No class due to Final Exam</p> <p>Abstract paper submission (2 columns * 2 page)</p>

Assignment and Project

Simulated robots will be set up using a Google Colab and controlled by students. The setup will be used for the assignments. The robots can be used for the final project if you want.

Each assignment is to be completed by you alone. You should not share the code with other students or other people. Note that this class does **not allow publishing any provided class content** in public (e.g., GitHub or blog).

The assignments should be submitted on the due date to receive full points. If not, there will be penalties following scale:

1. For each day thereof late, the instructor will reduce the final points you will receive by 20%.
*E.g.) If you submit it two days later, you will get 'total score * (100%-20%)** the number of dates'*
2. Late assignments over 2 days will receive zero points as a failure.

Note that the following policies

- **No E-MAIL submission** (Any E-mail submissions will not be allowed.)
- **No submission exception** (Any submission mistakes will not be accepted.)
- **PDF format only** (Assignments must be in **PDF** format. Otherwise, there will be 20% penalty.)
- **Code attachment** (Missing code attachment will get 50% penalty.)
- **English report only** (Any other language of reports will get zero scores.)

Paper presentation

In this course, we will take turn in reviewing and presenting a set of selected papers in this Robot Learning and Interaction area. Expect to present about one-two papers over the semester. You can select a paper from a google form and get a copy via the links above. Selected papers represent the recent advancements or fundamentals in the area.

The organization of your presentation depends on the type of paper. I'd like to recommend you address the following issues in your presentation:

1. Explain what the problem is
2. Explain why the problem is important/hard/valuable
3. Explain what the solution is
4. Explain the advantages/disadvantages of the approaches cited in paper
5. Describe good things about the paper
6. If there is an experiment, explain the setup
7. Describe major/minor comments

You should prepare about 45min of presentation. We will have about 15-20min for discussion. Then, we will have a quiz that the presenter provides. Your participation, quiz grade, and presentation will be recorded to determine your class participation score.

Course Evaluation

- Assignment (30%)
- Paper presentation + quiz (30%)
- Mid-term + Final presentations + abstract paper (30%)
- Class participation (10%)

Attendance

- This course requires attendance, which will be counted for grading your class participation score.
- If you miss three times of classes, you will lower your grade.
- When you cannot attend any lecture, if there is any online lecture, the recorded video will be posted via KLMS so you will be able to learn the class material and submit any quiz.

References

- Paper readings will be linked to specific topics in the schedule table.

<https://arxiv.org/pdf/2404.09080>

<https://ieeexplore.ieee.org/document/10933541?denied=>