

Georgia Institute of Technology / School of Interactive Computing
CS 7648 Interactive Robot Learning

1. Course Information

Semester:	Fall 2026
Academic Year:	2026 - 2027
Instructor:	Associate Professor Matthew Gombolay Matthew.Gombolay@cc.gatech.edu
Credit:	3-0-3
Pre-requisites:	None
Location:	TBD
Time:	TBD
Teaching Assistant:	TBD

2. General Course Description

2.1 Description

Humans can easily learn a new task by observing a demonstrator, asking a few questions, and gaining experience from a brief period of trial, error, and refinement. The field of Learning from Demonstration (LfD) has sought to endow robots with this human-like ability to learn through observation and interaction. In this course, we will cover challenges such as how to 1) learn from novice users (i.e., noisy, sparse demonstrations), 2) intelligently query the demonstrator for additional information (i.e., active learning), 3) compose hierarchical models to learn both high-level tasks and low-level motion primitives, and 4) evaluate the design of LfD systems via human-subject experimentation. The course culminates in a final research project in which students develop and demonstrate their own LfD technique, which will enable them to contribute to the deployment robotic technology in society.

2.2 Instructional Format

This course will be conducted primarily in a **chalk-talk format**. Rather than slide-based lectures, concepts will be developed interactively on the board, with an emphasis on deriving key ideas from first principles. Students are expected to actively participate in constructing models, questioning assumptions, and connecting concepts across papers. Students may be called upon to contribute to the development of ideas during class. Emphasis will be placed on formalizing problems through objective functions, probabilistic models, and algorithmic structure. The goal is depth of understanding rather than breadth of coverage.

2.3 Targeted Students

This course is designed for graduate students with a background in machine learning and reinforcement learning who are interested in interactive robot learning and human-robot interaction. Students should be comfortable with concepts from probability, optimization, and reinforcement learning (e.g., policies, value functions, and reward design), but prior exposure to human-robot interaction or learning from demonstration is not required.

2.4 Objectives and Expected Outcomes

This course aims to introduce the basic principles and techniques of LfD. By the end of the course, the students should be able to:

1. Articulate the value of LfD relative to more traditional forms of programming robots,
2. Identify and discuss key results and open problems from prior work in LfD,
3. Apply a variety of machine learning techniques to enable robot LfD,
4. Design and conduct an experiment to train and validate the design of their own LfD system.
5. Derive, analyze, and critique learning algorithms from first principles in an interactive setting.

3. Course Materials

3.1 Required Textbooks

- Chernova, S., & Thomaz, A. L. (2014). Robot learning from human teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(3), 1-121. ISBN: 978-3031004421.
- Russell, S.J. and Norvig, P., 2016. *Artificial Intelligence: A Modern Approach*. Malaysia; Pearson Education Limited. ISBN: 978-1292153964.

3.2 Optional Textbook Resources

- Mitchell, T. M. (1997). Machine learning. *McGraw Hill*. ISBN: 978-0071154673
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT press. ISBN: 978-0262193986.
- Montgomery, D. C. (2017). *Design and analysis of experiments*. John Wiley & Sons. ISBN: 978-1119113478.

4. Grading

4.1 Grade Components

Grades will reflect active participation, clarity of reasoning, correctness of derivations, and the ability to connect theory to practice. The final grade is composed as follows:

• Attendance and Participation	(individual)	10%
• Problem Sets (PSets)	(individual)	20%
• Paper Critiques	(individual)	10%
• Laboratory Activity	(team grade)	5%
• Midterm	(individual)	10%
Student-facilitated Seminar		
• Leading	(team grade)	10%
• Audience participation	(individual)	5%
Term Project		
• Proposal	(team grade)	4%
• Proposal Revision	(team grade)	4%
• Project update	(team grade)	12%
• Report/presentation	(team grade)	20%

4.1 Additional Criteria for Successful Completion of the Course

To successfully complete this course, students must complete all major course components, including the term project and student-facilitated seminar. Demonstrated understanding of core concepts, active participation, and timely submission of assignments are required to earn a passing grade.

5. Assignments and Coursework

5.1 PSets

During the first half of the semester, students will complete a series of individual Problem Sets (PSets). These assignments focus on developing a rigorous understanding of core concepts through derivations, analysis, and short technical exercises. Students are encouraged to discuss ideas with peers, but each student must independently write and submit their own solutions. Emphasis will be placed on clarity of reasoning, correctness of derivations, and the ability to justify modeling and algorithmic choices.

5.2 Paper Critiques

During the second half of the semester, as the course transitions to student-facilitated seminar sessions, students will complete short written critiques of the assigned papers. For each seminar session, all students will submit a brief (approximately 300 words) technical critique of the paper prior to the start of the class session in which the paper is discussed. These critiques are intended to prepare students to actively engage in discussion and evaluate research contributions critically. Each critique must include:

- Core Idea (1–2 sentences): A concise statement of the paper’s main contribution
- Critical Analysis: Identification and analysis of at least one key assumption, limitation, or potential failure mode
- Extension or Question: A proposed extension, alternative approach, or technically grounded question

At least one component of the critique must reference a specific model, objective function, or algorithmic element from the paper. Critiques will be evaluated based on:

- Clarity of Core Idea: Accurate and concise identification of the paper’s main contribution
- Depth of Analysis: Ability to identify and reason about key assumptions, limitations, or failure modes
- Technical Engagement: Use of concrete concepts (e.g., models, objectives, or algorithms) rather than high-level description
- Original Insight: Thoughtful extension, question, or alternative perspective

Critiques that primarily summarize the paper without analysis will not receive full credit.

5.3 Laboratory Activity

Students will complete a hands-on laboratory activity using a tabletop robotic platform (e.g., Kinova Jaco) to implement and evaluate a simple learning-from-demonstration or feedback-based pipeline. The goal of this activity is not to achieve high performance, but to expose the practical challenges of deploying interactive robot learning methods in physical systems. Students will work in their project groups to:

- implement a basic learning approach from demonstration or feedback
- collect and use data from real robot interactions
- evaluate the resulting behavior and identify failure modes

Students should expect the system to perform imperfectly. A central objective of the activity is to analyze the gap between theoretical methods and real-world performance, including sources of noise, variability in demonstrations, and limitations of the learning approach.

Students will submit a brief (maximum one page) group write-up summarizing their experience. The write-up should:

- briefly describe the approach implemented and data collected
- identify and analyze key failure modes observed during execution
- connect these observations to concepts from the course (e.g., distribution shift, noise in demonstrations, model mismatch, or optimization issues)

At least one component of the analysis must reference a specific modeling assumption or algorithmic element. Emphasis will be placed on depth of analysis rather than system performance.

Note: A sign-up system will be provided to schedule time with the robot platform.

5.4 Student-Facilitated Seminar

Students will be assigned to groups of three (3) by the instructor. Each group will also be assigned a seminar date and associated paper from the topical outline (below) and will be responsible for facilitating a class session on that paper. Each group should plan for approximately 75 minutes total, consisting of roughly 40 minutes of structured content development and 35 minutes of guided discussion.

All students are expected to be present in person for their assigned session. If a student is unable to attend, they must obtain an accommodation in advance through the appropriate Institute processes.

Sessions will be conducted in a **chalk-talk format**, in which students lead the class in reconstructing the key ideas, assumptions, and methods of the paper step-by-step. Rather than relying on slides, students are expected to develop the core concepts on the board, engage the class through questions and intermediate predictions, and guide critical discussion of the paper's contributions, limitations, and implications.

Student-led sessions will be evaluated based on:

- *Conceptual Development* – Clarity and correctness in building the core ideas, models, or algorithms from first principles
- *Structure and Pacing* – Logical progression of ideas and effective use of time
- *Engagement* – Ability to actively involve the class through questions, predictions, and discussion
- *Technical Depth* – Appropriate level of rigor in presenting assumptions, derivations, and limitations
- *Critical Insight* – Quality of analysis regarding the paper's contributions, weaknesses, and broader implications

Sessions that rely primarily on slide-based presentation rather than interactive development of ideas will not receive full credit.

5.5 Term Project

Students will conduct a group (2-3 students) research project based on the topics in this course. Students are encouraged to propose projects relevant to their own research to bring their unique perspectives. However, the course project must be something the student would not have done during their normal course of research and coursework. The term project consists of:

- Proposal – Two-page proposal detailing the project motivation, description, data, plan to collect the data, expected outcome, identification of benchmark, and timeline. The proposal should address the Heilmeier Catechism (<https://www.darpa.mil/work-with-us/heilmeier-catechism>). The instruction team will provide feedback on the proposal, which should be incorporated into the project update and final deliverable.
- Proposal Revision – Teams will revise their two-page proposal based upon the feedback from the instruction team. The revision will be graded separately (4%) and evaluated based on how effectively teams address feedback, refine their problem formulation, and improve the clarity and feasibility of their approach.
- IRB Application – Training and evaluation of LfD requires working with human-subjects. If you want to publish with the data you collect during this project, you *must* receive Internal Review Board (IRB) approval. Each group is encouraged to submit one application. The IRB application is ungraded and only required for students wishing to disseminate the results of their course projects in the broader academic community (e.g., a workshop, conference, or journal paper). **However,**
 - **project teams who submit an IRB application will receive +1 on their final grade, and**
 - **project teams who subsequently get approval for their application, which may require revisions and iteration with the IRB, will get an additional +1 (total of +2) on their final grade.**

Please note: Dr. Gombolay is required to certify any IRB application, and he reserves the right to reject an application if it proposes unethical experimentation or is otherwise deficient.

- Project Update – A one-page summary and a 10-minute, in-class presentation detailing the progress to date on executing the project proposal. Presenters' peers will provide feedback. The final project presentation grade will be influenced by how well students incorporate helpful, reasonable feedback into their projects.
- Project Presentation – An 8-page conference-style paper and a 10-minute, in-class presentation detailing students' contributions for their research project. The paper should have an abstract, introduction, related works, methods, results, discussion, and conclusion. The results section must include a benchmark that the students applied to their data set.

Note: term project presentations are expected to be in-person, and remote participation will not be possible without an excused absence.

Three projects in this course have resulted in arxiv/workshop publications, which are provided here as examples of relatively good projects:

Luo, Q., Wu, J. and Gombolay, M., 2020. A Generalized Robotic Handwriting Learning System based on Dynamic Movement Primitives (DMPs). arXiv preprint arXiv:2012.03898.

Zuo, M., Schick, L., Gombolay, M. and Gopalan, N., 2022. Learning Efficient Exploration through Human Seeded Rapidly-exploring Random Trees. arXiv preprint arXiv:2203.12774.

Fernandez, D.J., Grande, G.N., Sean, C.Y., Gopalan, N. and Gombolay, M., 2022. Participatory Design and End-User Programming for Human-Robot Interaction. In *Proceedings of the IEEE/ACM International Conference on Human-Robot Interaction Workshop on Participatory Design and End-User Programming for Human-Robot Interaction*. Available from https://bpb-us-w2.wpmucdn.com/sites.gatech.edu/dist/d/958/files/2022/03/Interactive_Learning_HRI_Workshop.pdf

6. Course Policies

6.1 Attendance/Participation

100% attendance is expected to provide students with the best opportunity to learn and contribute to the learning of their peers through constructive interaction. In-class participation will be assessed through engagement with the live development of ideas during lecture (e.g., responding to questions, contributing to derivations, and brief concept checks). Short, low-stakes check-ins may be used to support engagement. Excused absences are permitted and require official institute documentation, as described here: <https://registrar.gatech.edu/info/institute-approved-absence-form-for-students>. If you require an accommodation that may affect attendance or participation, please notify the instructor in advance and, where appropriate, coordinate with the Office of Disability Services.

6.2 Late Policy

Late assignments, except for the final project, will be accepted with 1 letter grade off per 1 day, rounding up (e.g., 1 second and 23 hours 59 minutes 59 seconds each result in 1 letter grade off). No late assignments will be accepted for the final project presentation. In the case of an excused absence, the student and teacher will work to arrange an extension under the guidelines of GaTech.

6.3 In-person Expectations

This course is conducted in person and relies on active, real-time participation. Students are expected to attend class regularly and be prepared to engage in discussion, contribute to the development of ideas, and participate in in-class activities. Student-facilitated seminar sessions and project presentations must be completed in person. If a student is unable to attend a required in-person activity, they must notify the instructor in advance and, where appropriate, obtain accommodation through Institute processes.

7. Schedule and Topics

Week	Date	Topic	Reading Assignments Due	Graded Assignments Due
01	08/25	Introduction to Interactive Robot Learning	C. & T. – Chapters 1-3 Russell & Norvig Chapter 21 Thomaz, A., & Breazeal, C. (2008) Meltzoff, A. N. (2005)	
	08/27	Behavior Cloning (BC) and Dataset Aggregation (DAgger) Sequence Modeling	Ross, et al. (2011) Spencer et al. (2020) Chen et al. (NeurIPS-21)	
02	09/01	Inverse Reinforcement Learning – Max Margin	Abbeel, P., & Ng, A. Y. (2004)	
	09/03	Inverse Reinforcement Learning – Max Entropy	Ziebart, B.D. et al. (2008)	Pset1 Due: 09/04 @ 23:59 Eastern

03	09/08	Deep, Generative, and Adversarial LfD	Fu et al. (2017) Ho & Ermon (2016)	
	09/10	Preference Learning	Christiano et al., 2017 Hejna et al. (2023)	Pset2 Due: 09/11 @ 23:59 Eastern
04	09/15	Unifying View: Learning from Different Signals	Levine (2018)	
	09/17	Learning from feedback and advice	C. & T. – Chapter 6; Knox & Stone (2009) Celemin & Ruiz-del-Solar (2015)	Pset3 Due: 09/18 @ 23:59 Eastern
05	09/22	Active Learning	Chernova & Veloso (2009) Schrum & Gombolay (2022)	
	09/24	Human Modeling	Nikolaidis et al. (2015) Xue et al. (2025)	Pset4 Due: 09/25 @ 23:59 Eastern
06	09/29	Language as Conditioning	Silva et al. (2022) Tambwekar et al. (2023)	
	10/01	Evaluating LfD	C. & T. – Chapter 7	Pset5 Due: 10/02 @ 23:59 Eastern
07	10/06	Low-level Skills (Motion Primitives)	Saveriano et al. (2023) C. & T. – Chapter 4	
	10/08	High-level Skills (Hierarchical Task Networks)	Chen et al. (CoRL-21) C. & T. – Chapter 5	Pset6 Due: 10/09 @ 23:59 Eastern
08	10/13	Future Directions & Project Kickoff!	C. & T. – Chapter 8	
	10/15	<i>In-class Midterm</i>		

Student Lectures Begin:

09	10/20	LfD via Skill Trees	Konidaris, G., et al. (2012)	
	10/22	Probabilistic Movement Primitives	Paraschos et al. (2018)	Project Proposals Due: 10/23 @ 23:59 Eastern
10	10/27	LfD in Table Tennis	Mülling et al. (2023)	
	10/29	Proposal workshop day		
11	11/03	Active Learning via Imitation Learning	Bullard et al., (2019)	
	11/05	LfD In the Wild	Behbahani et al. (2019)	Proposal Revisions Due: 11/07 @ 23:59 EST
12	11/10	Grounded Semantic Reasoning for Robotic Interaction with Real-World Objects	Liu et al. (2022)	
	11/12	Do As I Can, Not As I Say	Ahn et al. (2022)	
13	11/17	RT-2	Zitkovich et al. (2023)	
	11/19	No, to the Right	Cui et al. (2023)	
14	11/24	Robots Teaching Robots to be Better Robot Teachers	Schrum et al. (2022) Schrum et al. (2023)	Project updates Due: 11/25 @ 23:59 Eastern
	11/26	Class Holiday		
15	12/01	Viola	Zhu et al. (2023)	
	12/03	Safety in LfD	Yang et al. (2024)	

16	12/08	Fictitious Co-Play	Strouse et al. (2021)	Laboratory Activity Due: 12/08 @ 23:59 Eastern
17	TBD	Project Presentations		Project Presentations Due: to Prof. Gombolay by TBD.

8. Institutional Policies

8.1 Accommodations

If you are a student with learning needs that require special accommodation, contact the Office of Disability Services at (404) 894-2563 or <http://disabilityservices.gatech.edu/>, as soon as possible, to discuss your needs and to obtain an accommodations letter. Please e-mail me as soon as possible to set up a time to discuss your learning needs.

8.2 Academic Integrity

Georgia Tech aims to cultivate a community based on trust, academic integrity, and honor. Students are expected to act according to the highest ethical standards. For information on Georgia Tech's Academic Honor Code, please visit <http://www.catalog.gatech.edu/policies/honor-code/%20or%20http://www.catalog.gatech.edu/rules/18/>. Any student suspected of cheating or plagiarizing will be reported to the Office of Student Integrity.

Students are not allowed to use any language model (e.g., CodeLlama, ChatGPT, etc.) to help you with any graded component of this course. The instruction team reserves the right to use automated tools to evaluate student work.

8.3 Student-Faculty Expectations Agreement

At GaTech, we believe that it is important to strive for an atmosphere of mutual respect, acknowledgement, and responsibility between faculty members and the student body. See <http://www.catalog.gatech.edu/rules/22/> for an articulation of **expectations** you can have of me and I have of you. Respect for knowledge, hard work, and cordial interactions will help build the environment we seek.

8.4 Statement of Intent for Inclusivity

As members of the Georgia Tech community, we are committed to creating a learning environment in which all students feel safe and included. Because we are individuals with varying needs, we are reliant on your feedback to achieve this goal. To that end, we invite you to enter into dialogue with us about the things we can stop, start, and continue doing to make our classroom an environment in which every student feels valued and can engage actively in our learning community.

8.5 Amendments

Your instructors reserve the right to make changes as severe weather and other factors necessitate. Any changes will be accompanied by notice from the instructors.

9. Communication and Office Hours

9.1 Office Hours

We will be using office hours as a mechanism to facilitate live interactions between instructors/TA's and enrolled students. All hours below are Eastern Time.

Instructor Office Hours: TBD
TA Office Hours: TBD

9.2 Communication

- Course communication will be conducted primarily through Piazza. Students are expected to use Piazza for all course-related questions, including questions about assignments, course content, and logistics. Posting questions publicly is strongly encouraged so that all students can benefit from the discussion.
- Posts on Piazza may be made anonymously to classmates; however, student identities are visible to the instructional team. This allows for open participation while maintaining accountability.
- The teaching assistants will serve as the first point of contact on Piazza and will respond to most questions. The instructor will monitor discussions and intervene as needed.
- Students are encouraged to search existing posts before asking new questions and to contribute to discussions when appropriate.
- Questions sent via email about course content may be redirected to Piazza.
- For private matters (e.g., personal or administrative concerns), students may contact the instructional team directly.

10. References

Note: Some references below are supplementary and are not required reading. Required reading assignments are listed in Section 7.

- Abbeel, P., & Ng, A. Y. (2004, July). Apprenticeship learning via inverse reinforcement learning. In *Proceedings of International Conference on Machine Learning*.
- Ahn, M., Brohan, A., Brown, N., Chebotar, Y., Cortes, O., David, B., Finn, C., Fu, C., Gopalakrishnan, K., Hausman, K. and Herzog, A., 2022. Do as I can, not as I say: Grounding language in robotic affordances. arXiv preprint arXiv:2204.01691.
- Behbahani, F., Shiarlis, K., Chen, X., Kurin, V., Kasewa, S., Stirbu, C., Gomes, J., Paul, S., Oliehoek, F.A., Messias, J. and Whiteson, S., 2019, May. Learning from demonstration in the wild. In 2019 International Conference on Robotics and Automation (ICRA) (pp. 775-781). IEEE.
- Bullard, K., Schroecker, Y. and Chernova, S., 2019. Active learning within constrained environments through imitation of an expert questioner. arXiv preprint arXiv:1907.00921.
- Celemin, C. and Ruiz-del-Solar, J., 2015. COACH: learning continuous actions from corrective advice communicated by humans. In *Proc. International Conference on Advanced Robotics (ICAR)* (pp. 581-586).
- Chen, K., Srikanth, N.S., Kent, D., Ravichandar, H. and Chernova, S., 2021. Learning hierarchical task networks with preferences from unannotated demonstrations. In *Conference on Robot Learning* (pp. 1572-1581). PMLR.
- Chen, L., Paleja, R., Ghuy, M. and Gombolay, M., 2020. Joint goal and strategy inference across heterogeneous demonstrators via reward network distillation. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 659-668).
- Chen, L., Paleja, R. and Gombolay, M., 2020. Learning from suboptimal demonstration via self-supervised reward regression. arXiv preprint arXiv:2010.11723.
- Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A. and Mordatch, I., 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34, pp.15084-15097.
- Chernova, S., & Veloso, M. (2009). Interactive policy learning through confidence-based autonomy. *Journal of Artificial Intelligence Research*, 34, 1-25.
- Christiano, P.F., Leike, J., Brown, T., Martic, M., Legg, S. and Amodei, D., 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Cui, Y., Karamcheti, S., Palleti, R., Shivakumar, N., Liang, P. and Sadigh, D., 2023, March. No, to the Right: Online Language Corrections for Robotic Manipulation via Shared Autonomy. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 93-101).
- Duan, Y., Andrychowicz, M., Stadie, B., Ho, O.J., Schneider, J., Sutskever, I., Abbeel, P. and Zaremba, W., 2017. One-shot imitation learning. In *Advances in neural information processing systems* (pp. 1087-1098).
- Fu, J., Luo, K. and Levine, S., 2017. Learning robust rewards with adversarial inverse reinforcement learning. arXiv preprint arXiv:1710.11248.

- Gombolay, M., Jensen, R. and Son, S.-H., 2017. Machine learning techniques for analyzing training behavior in serious gaming. *IEEE Transactions on Computational Intelligence and AI in Games*.
- Hejna, J., Rafailov, R., Sikchi, H., Finn, C., Niekum, S., Knox, W.B. and Sadigh, D., 2023. Contrastive Preference Learning: Learning from Human Feedback without RL. arXiv preprint arXiv:2310.13639.
- Ho, J. and Ermon, S., 2016, December. Generative adversarial imitation learning. In Proceedings of the 30th International Conference on Neural Information Processing Systems (pp. 4572-4580).
- Knox, W. B., & Stone, P. (2009). Interactively shaping agents via human reinforcement: The TAMER framework. In *Proceedings of the fifth international conference on Knowledge capture* (pp. 9-16).
- Konidaris, G., Kuindersma, S., Grupen, R., & Barto, A. (2012). Robot learning from demonstration by constructing skill trees. *The International Journal of Robotics Research*, 31(3), 360-375.
- Levine, S., 2018. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*.
- Liu, W., Paxton, C., Hermans, T. and Fox, D., 2022, May. Structformer: Learning spatial structure for language-guided semantic rearrangement of novel objects. In 2022 International Conference on Robotics and Automation (ICRA) (pp. 6322-6329). IEEE.
- Meltzoff, A. N. (2005). "Imitation and other minds: The "Like Me" hypothesis." In S. Hurley and N. Chater (Eds.), *Perspectives on Imitation: From Neuroscience to Social Science* (Vol 2, pg 55-77), Cambridge, MA. MIT Press.
- Mülling, K., Kober, J., Kroemer, O. and Peters, J., 2013. Learning to select and generalize striking movements in robot table tennis. *The International Journal of Robotics Research*, 32(3), pp.263-279.
- Murali, A., Liu, W., Marino, K., Chernova, S. and Gupta, A., 2020. Same object, different grasps: Data and semantic knowledge for task-oriented grasping. arXiv preprint arXiv:2011.06431.
- Nikolaidis, S., Ramakrishnan, R., Gu, K. and Shah, J., 2015, March. Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In 2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI) (pp. 189-196). IEEE.
- Nicolescu, M. N., & Mataric, M. J. (2003, July). Natural methods for robot task learning: Instructive demonstrations, generalization and practice. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems* (pp. 241-248). ACM.
- Paleja, R., Silva, A., Chen, L. and Gombolay, M., 2020. Interpretable and Personalized Apprenticeship Scheduling: Learning Interpretable Scheduling Policies from Heterogeneous User Demonstrations. *Advances in Neural Information Processing Systems*, 33.
- Paraschos, A., Daniel, C., Peters, J. and Neumann, G., 2018. Using probabilistic movement primitives in robotics. *Autonomous Robots*, 42, pp.529-551.
- Ramachandran, D., and Eyal A. (2009). Bayesian inverse reinforcement learning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)* (pp 2586-2591).
- Ross, S., Gordon, G. and Bagnell, D. (2011). A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics* (pp. 627-635).
- Saveriano, M., Abu-Dakka, F.J., Kramberger, A. and Peternel, L., 2023. Dynamic movement primitives in robotics: A tutorial survey. *The International Journal of Robotics Research*, 42(13), pp.1133-1184.
- Schrump, M., Connolly, M.J., Cole, E., Ghetiya, M., Gross, R. and Gombolay, M.C., 2022. Meta-Active Learning in Probabilistically Safe Optimization. *IEEE Robotics and Automation Letters*, 7(4), pp.10713-10720.
- Schrump, M.L., Hedlund-Botti, E., Moorman, N. and Gombolay, M.C., 2022, March. Mind meld: Personalized meta-learning for robot-centric imitation learning. In 2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI) (pp. 157-165). IEEE.
- Schrump, M.L., Hedlund-Botti, E. and Gombolay, M., 2023, March. Reciprocal mind meld: Improving learning from demonstration via personalized, reciprocal teaching. In *Conference on Robot Learning* (pp. 956-966). PMLR.
- Silva, A., Moorman, N., Silva, W., Zaidi, Z., Gopalan, N. and Gombolay, M., 2021. LanCon-Learn: Learning With Language to Enable Generalization in Multi-Task Manipulation. *IEEE Robotics and Automation Letters*, 7(2), pp.1635-1642.

- Spencer, J., Choudhury, S., Barnes, M., Schmittle, M., Chiang, M., Ramadge, P. and Srinivasa, S. (2020). Learning from Interventions: Human-robot interaction as both explicit and implicit feedback. In *Proceedings of Robotics: Science and Systems*.
- Strouse, D.J., McKee, K., Botvinick, M., Hughes, E. and Everett, R., 2021. Collaborating with humans without human data. *Advances in neural information processing systems*, 34, pp.14502-14515.
- Sun, W., Venkatraman, A., Gordon, G. J., Boots, B., & Bagnell, J. A. (2017, July). Deeply AggreVaTeD: Differentiable Imitation Learning for Sequential Prediction. In *International Conference on Machine Learning* (pp. 3309-3318).
- Thomaz, A. L., & Breazeal, C. (2008). Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, 172(6-7), 716-737.
- Xue, C., Chen, L. and Gombolay, M., 2025. Better Than Diverse Demonstrators: Reward Decomposition from Suboptimal and Heterogeneous Demonstrations. *IEEE Robotics and Automation Letters*.
- Yang, Y., Chen, L., Zaidi, Z., van Waveren, S., Krishna, A. and Gombolay, M., 2024, March. Enhancing safety in learning from demonstration algorithms via control barrier function shielding. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 820-829).
- Yu, T., Finn, C., Xie, A., Dasari, S., Zhang, T., Abbeel, P. and Levine, S., 2018. One-shot imitation from observing humans via domain-adaptive meta-learning. arXiv preprint arXiv:1802.01557.
- Ziebart, B.D., Maas, A., Bagnell, J.A. and Dey, A.K., 2008. Maximum entropy inverse reinforcement learning. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*.
- Zhu, Y., Joshi, A., Stone, P. and Zhu, Y., 2023, March. VIOLA: Object-Centric Imitation Learning for Vision-Based Robot Manipulation. In *Conference on Robot Learning* (pp. 1199-1210). PMLR.
- Zitkovich, B., Yu, T., Xu, S., Xu, P., Xiao, T., Xia, F., Wu, J., Wohlhart, P., Welker, S., Wahid, A. and Vuong, Q., 2023, December. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning* (pp. 2165-2183). PMLR.