huaol@mit.edu

Huao Li

www.huao-li.com

Teaching Statement

My commitment to education is rooted partly in my family history: my grandfather spent his career teaching in a rural, underserved region of China, and his example made education feel to me not only like knowledge transfer, but also a form of service. I have also been fortunate to learn from teachers and advisors whose encouragement and intellectual freedom made my own path to a Ph.D. possible and profoundly shaped my values and career choices. Their example taught me to see education as supporting the growth of others, a purpose I hope to make central to my own career.

This goal takes on new meaning in the era of generative AI. Students are preparing for careers in which AI will increasingly participate in their work, yet many courses still assume that they will solve problems independently from beginning to end. Learning foundational concepts and skills remains essential, but critical thinking and the ability to collaborate with AI and other humans are becoming increasingly important. My goal is to help students develop these abilities through authentic assignments that require them to evaluate AI contributions, question assumptions, and work with others to solve problems.

Teaching Philosophy

Learning through authentic practice. I use backward design to connect learning objectives with the tasks students undertake and the evidence they produce. Rather than beginning with the material I want to cover, I start from what students should be able to understand or accomplish and design activities through which they can develop and demonstrate those abilities.

For example, asking students to implement a standard reinforcement-learning algorithm from scratch no longer provides sufficient evidence of understanding when AI can generate a plausible implementation. In a future RL course, I would allow AI assistance in a project that asks students to develop an agent for a task with competing objectives and evaluate it under conditions that differ from training. Students would need to justify their reward design, investigate failure cases, and revise their approach based on evidence. Assessment would consider both system performance and the reasoning behind their decisions: what they delegated to AI, how they checked its contributions, and why they accepted or changed its suggestions. Targeted exercises would still develop foundational skills, while the project would require students to apply those skills in a realistic workflow.

MIT's Kaufman Teaching Certificate Program (KTCP) helped me understand that this focus on learning outcomes also requires flexibility during instruction. In a lecture on value alignment, I asked students how an autonomous vehicle should resolve a conflict between human values. I had expected a few discrete positions; instead, they raised qualifications that challenged the assumptions behind the question. I extended the discussion because it directly served the learning objective: understanding why alignment is difficult when people do not share a single value system. This experience changed how I understood student-centered teaching. Following a lecture plan is not itself evidence of learning, and productive disagreement can advance an objective more effectively than another explanation from the instructor.

Teaching across backgrounds. As AI tools create new opportunities to engage with computing, students may enter computational classes with very different programming experience and disciplinary knowledge. I address this variation through instructional scaffolding: making the pathway into a challenging task accessible without lowering the level students can ultimately reach.

As a teaching assistant for Human-Robot Interaction at the University of Pittsburgh, I redesigned a reinforcement learning project around this principle. Students received a working task environment and example implementation, then implemented algorithms discussed in class and evaluated them on a shared benchmark. This structure reduced the initial programming burden while preserving the task of implementing and comparing learning methods. For advanced students, I included an open-ended challenge to improve performance in a highway environment through approaches such as reward shaping or representation learning. A leaderboard gave them a shared reference for comparing their results.

Having moved across different educational systems and departments, I remember the difficulty of entering a field whose assumptions others take for granted. In my courses, I would use short preparation surveys to identify differences in experience and structured small-group discussions to let students develop

questions before speaking to the full class. My goal is to make these differences resources for learning while ensuring that every student develops the technical understanding required by the course.

From theory to real-world impact. My interdisciplinary training shapes what I hope students learn to build and how they evaluate it. Computational performance matters, but it does not by itself establish whether a system addresses human needs. Students should learn to connect technical decisions with evidence about the people who use or are affected by a system.

During KTCP, I designed a graduate course, *Introduction to Human-Centered AI*. Building on this design, I would organize the course around human modeling, learning from human feedback, interpretability, human-AI teaming, and evaluation of societal impact. Students would examine how assumptions about people enter computational methods and how those assumptions can be tested. The final project allows students to pursue a design and evaluation study of an AI product, a computational model of a human cognitive process, or an analysis of an AI system's societal impact. Each track would require a clear question, a justified method, and evidence supporting its conclusions. This flexibility allows students to draw on their own training while holding all projects to a common standard of rigor.

Group projects would also give students practice in communication and project management. To facilitate these learning outcomes, I would form teams with complementary experience and establish shared milestones and explicit responsibilities. Rotating responsibilities would help prevent students with less experience from being confined to non-essential work. Peer feedback and individual introspection of project decisions would support accountability alongside the shared outcome.

Mentoring Experience

Across my research training and postdoctoral work, I have mentored ten students at the high-school, undergraduate, and master's levels. Projects I have mentored have resulted in seven publications, including work at AAMAS, AIES and IEEE Transactions. I have also learned from mentees whose technical strengths and perspectives differ from my own. These experiences have convinced me that mentoring is an extension of teaching: the goal is not simply short-term research output, but helping students learn to frame problems, respond to failure, seek feedback, and eventually become independent researchers.

I treat mentees as collaborators whose ideas deserve real consideration, and I leave room for those ideas to develop. One high-school student I mentored through MIT's Research Science Institute arrived interested in AI with almost no background in reinforcement learning. I began with foundational readings, tutorials, and a narrowly scoped question, then handed over more of the project as he became more capable and confident. When his early experiments failed to support our hypothesis, we worked through how to diagnose the result, what else could explain it, and which experiment would distinguish the possibilities. By the end of the summer, he was analyzing results on his own and proposing follow-up studies. The amount of structure I provide changes with the student and the stage of the project, and I deliberately reduce that structure as their confidence and independence grow.

My second mentoring principle is to build a community in which students can exchange ideas and seek help without feeling that uncertainty signals a lack of ability. At MIT, I have helped organize departmental postdoc meetings and participated in peer-support activities within the Engineering Excellence postdoctoral program. In my own group, I would complement individual advising with discussions of work in progress, including unsuccessful experiments, and opportunities for students to give one another feedback. These practices would make research difficulties shared opportunities for learning while preserving each student's intellectual independence.

Course Plan

At the undergraduate and graduate levels, I am prepared to teach introduction to artificial intelligence, machine learning, reinforcement learning, and human-computer interaction. I would also bring *Introduction to Human-Centered AI*, the new course described above. Across these courses, I would connect technical foundations with authentic opportunities to build and evaluate systems with AI assistance. My goal is for students to leave with both a strong understanding of the subject and the ability to direct, evaluate, and take responsibility for work completed with increasingly capable artificial agents.