Teaching Statement: Teaching to Achieve Meta-learning

Paul Pu Liang, Machine Learning Department, School of Computer Science, Carnegie Mellon University

My teaching experience has centered around courses in artificial intelligence (AI), specifically in machine learning (ML) and deep learning for real-world multimodal data such as language, images, video, audio, and sensors. **I have co-led the revamp and lecturing of the core AI course** on Multimodal ML (CMU 11-777 in Fall 2019, Fall 2020, Fall 2022, Fall 2023) and **co-created two special topics courses** on Advanced Multimodal ML (CMU 11-877 in Spring 2022, Spring 2023) and Artificial Social Intelligence (CMU 11-866 in Spring 2023). **I also served as a teaching assistant 4 times** for Probabilistic Graphical Models, Advanced ML, Intro to ML, and Computer Systems. Finally, **I have led and advised more than 50 students in directed research**, with 15 from underrepresented groups: 13 of them have gone to top PhD programs at CMU, UC Berkeley, MIT, USC, UW, Princeton, University of Michigan, Virginia Tech, Yale, KAIST, and the University of Tokyo, 4 of them received CRA undergrad research awards, and 4 of them received summer research fellowships. For my teaching contributions, I received the **Alan J. Perlis Graduate Student Teaching Award** in 2023 and the **Teaching Assistant Award** in 2019. As faculty, I am excited to teach core courses in machine learning, deep learning, artificial intelligence, natural language processing, data science, and statistics, as well as special topics courses in multimodal machine learning.

The recent wave of AI has exploded since 2010, with deep learning bringing disruptive advances in methods and applications. This has made teaching challenging since entire ideas that did not exist merely years before now proliferate all aspects of research and development. How can we ensure that our students are learning the most up-to-date material for tangible real-world impact? Are there better ways of teaching that encourage students to come up with the next wave of innovation themselves? These questions have raised unique opportunities to explore new teaching methodologies, and I summarize my philosophy through the term *meta-learning*.

Learning vs meta-learning: In his seminal book on Machine Learning, Tom Mitchell defines learning as:

One learns from experience E with respect to some tasks T and performance measure P, if their performance on tasks T, as measured by P, improves with experience E.

While this is a simplified model and does not capture the long-term positive externalities of teaching, there are still insights to be gained: students generally gain knowledge through lectures, recitations, and homework (experience E) on a specified set of topics predefined in a syllabus (tasks T) with the goal that they improve on some measurable score (performance measure P) with increased experience E. However, this may break down for graduate-level courses meant to tackle the most recent research topics. The new and fast-moving pace of this content (new tasks in T) makes it hard to develop pedagogically sound instructional and assessment material (experience E).

Meta-learning, as its name suggests, differs from standard learning with the goal of teaching students the process of *learning to learn*, rather than what to learn. In a similar vein as above, one could define meta-learning as:

One ***meta-learns*** *from experience E with respect to some tasks T and performance measure P, if their performance on* ***new tasks*** *from a similar distribution but not seen in T, as measured by P, improves with experience E.*

How can we structure teaching such that we achieve student meta-learning? Firstly, teaching fundamental concepts that permeate past, current, and future methods so that students can make significant contributions immediately and long into the future. Secondly, by encouraging curiosity to seek out open questions and accomplish open-ended research, with the goal of coming up with the next paradigm shift in ideas themselves.

1. Teaching concepts that permeate current and future methods

1.1 From fundamentals to applications. I co-taught Multimodal ML and completely revamped the course in Fall 2022 to highlight the fundamental principles and technical challenges inspiring the recent waves of applications, instead of focusing on the popular methods themselves. This revamp includes a completely new set of lecture notes, slides, reading materials, and project assignments to emphasize the fundamental ideas in mathematics, statistics, logical thinking, and programming. When giving lectures, I used illustrative figures to provide clear intuitions for deriving several classic multimodal algorithms, before emphasizing different implementations in the past, present, and what we expect in the future. I made it a point to be extra responsive in person and on Piazza to ensure everyone understood these fundamental concepts and students widely praised my lectures for deriving these basic theorems which they had previously seen but never completely understood. However, there were also students who asked why I was teaching these concepts when the most popular methods abstract away these details. Aware of the challenges in catering to students of diverse backgrounds and goals, I updated my content and delivery over the 3 iterations of multimodal ML to improve the balance between fundamentals and applications, high-level intuitions and low-level details, and modern advances and historical progressions.

1.2 Open discussion probes and broad readings. I have also co-created two special topics courses: Advanced Multimodal ML and Artificial Social Intelligence in Spring 2022 and 2023 to cover the most recent advances in these vibrant subfields of AI. However, we do not want students to 'overfit' to recent methods such as data collection, learning algorithms, and model architectures, since they will likely be outdated in just a few years' time. To help students gain more holistic perspectives on the current state-of-the-art and open directions, I designed a set of weekly discussion probes on open questions and provided an initial list of ten readings from different perspectives, time periods, and research areas. Students were tasked to read these papers, think critically about the questions, and scout for additional resources to share with the class. When leading group discussions, I prompt students to compare and contrast work from different perspectives while making sure that every student is comfortable voicing their thoughts through constructive discussion. Students widely praised this course format for teaching them how to read and critique research papers, broadening their knowledge of AI from different disciplines and time periods, and better preparing them to solve challenging open problems.

1.3 Broader impact of technology. Finally, I strongly believe in teaching students to reason about the broader impacts of technology on society, such as understanding AI robustness, fairness, interpretability, and accountability. While the specific method may shift over time, these real-world concerns will remain. I have created and delivered new lectures on the societal impact of AI in our courses and a guest lecture at another CMU course on human-AI interaction, while further inviting guest speakers who are experts on AI bias, fairness, and healthcare applications. Finally, I make sure every week's discussion includes at least 1 paper on the broader impact of technology and dedicate at least 15 minutes of the discussion to these issues. Based on these discussions, I also guide students in writing about the broader impact of their own research projects in their final reports. These methods encourage students to become well-rounded thinkers and practitioners from both technical and ethical perspectives.

2. Encouraging curiosity and enabling access to open-ended research

2.1 Emphasizing open-ended research. Perhaps the best way to prepare students for meta-learning is to encourage their curiosity to seek out new tasks by diving deep into an open question themselves. I have mentored more than 100 students for course research projects and 50 for directed research. I work closely with each student to design a holistic research plan, teach them research skills, and encourage their own research interests. Through regular meetings and hands-on co-working, I give guidance on identifying problems of impact, formulating research questions, searching for relevant literature, designing experiments, analyzing results, and communicating findings. Almost all students have co-authored publications at the top international conferences in ML and AI, with 15 of them writing first-author papers as a sign of exceptional research independence and maturity. I have shared my experiences of advising students as a member of the SCS PhD Dean's Advisory Committee Undergraduate Research Engagement Working Group by organizing research mixers and panels connecting prospective undergrads with grad student advisors to improve the participation, support, and experience of undergrad research.

2.2 Enabling broader access. I strongly believe that education should be a right and not a privilege, so I have dedicated efforts toward broader access to CS education. This past summer, I created and taught a week-long course on AI at the African Masters of Machine Intelligence program in Senegal, a program to train the next generation of researchers from Africa. I am now advising 7 students on building AI to comprehend Amharic language, on sentiment and emotion analysis for African speakers, and predicting energy grid usage and crop growth. For our CMU courses, I have made all lectures and materials available online. More broadly, I have taught tutorials at international conferences, and guest lectured at many universities and summer schools in North America, Europe, and Africa. They are regularly attended by thousands of participants across the world.

While I am proud of my efforts so far, I am constantly exploring new teaching ideas in collaboration with faculty and students. I look forward to teaching and mentoring the next generation of students by fostering fundamental inquiry, curiosity for research, and inclusive access.