

Philosophy

Shaping the next generation of leaders in technology and society is one of the primary inspirations behind my pursuit of an academic career. My teaching philosophy revolves around making learning engaging, accessible, and enjoyable. To achieve this, I plan to incorporate a thoughtful blend of:

- **Deep Exploration:** Leveraging the whiteboard to dive into topics with depth and clarity, fostering a “slow-learning” approach that builds a strong conceptual foundation.
- **Hands-On Engagement:** Encouraging “fast-learning” through interactive exploration and programming exercises that bring key concepts to life.
- **Collaborative Growth:** Facilitating “group learning” through team-based projects that promote peer collaboration, communication, and shared discovery.

This holistic approach aims to equip students not only with technical skills but also with the confidence and curiosity to lead and innovate.

Teaching Experience

Teaching at Facebook AI Research (FAIR). At FAIR, I contributed to the AI Residency Program, which supported promising undergraduates and researchers from non-ML fields (e.g., neuroscience) in launching their ML research careers. I volunteered to provide feedback on the program’s structure and taught lectures on topics like Meta-Learning. Gratifyingly, after one lecture, a student reached out to express their enthusiasm, sharing that my lecture had introduced them to exciting new techniques and inspired them to apply these methods to their own work.

Graduate Teaching Experience. During my master’s program at Virginia Tech, I served twice as a graduate teaching assistant for an undergraduate course on Computer Engineering (ECE 2504) that covered foundational topics from digital logic circuits to assembly-level programming. My responsibilities included grading, leading weekly lab sessions, and holding office hours. Interacting with students individually was particularly rewarding. It helped me appreciate how hands-on engagement enables students to form their own unique, individual perspectives on concepts.

During my second year of graduate school, I was honored with the “Best Discussion Participant Award” at the end of an “Advanced Computer Vision” course taught by Prof. Devi Parikh, where students led presentations on advanced topics in a class of 10-15 students. This award in many ways reflected my own individual growth as an aspiring researcher, helping me appreciate the power of group learning and discussion in fostering intellectual depth.

At Georgia Tech, I founded and ran a reading group called “UnArXiv” for over 1.5 years, focused on studying books and seminal papers in applied mathematics, optimization, and machine learning. The group’s unique rule – using only whiteboard explanations – encouraged deeper exploration of equations and concepts without the constraints of slides. A participant, Prithvijit Chattopadhyay (now a Research Scientist at NVIDIA) noted that UnArXiv discussions, particularly my ability to connect seemingly disparate ideas, continue to influence their approach to research even today.

Mentorship. At FAIR, I mentored two AI residents and three summer interns, co-mentored two students at NYU, and two others during graduate school. Remarkably, eight out of nine projects led to publications in top machine learning conferences. My mentoring philosophy emphasizes mentees’ motivations and tailors guidance to their unique strengths, inclinations and goals.

The feedback from my mentees reflects this approach. Yann Dubois (now a Stanford Ph.D. student) noted the balance I struck between giving him autonomy and engaging in thought-provoking discussions, as well as my genuine care for his well-being. Sirui Xie (UCLA Ph.D. student, now Research Scientist at Google) highlighted how I inspired his scientific and personal growth by fostering critical thinking, the art of experimental design, and helped him “break free from the authorities of language”. Similarly, Daksh Idnani (current CTO of his startup) praised my focus on developing systematic evaluation and effective communication skills, while Karan Desai (Michigan Ph.D. student, now Member of Technical Staff at an AI Startup) appreciated my commitment to foundational understanding and fundamental contributions over rushing for short-term results.

I encourage mentees to pursue meaningful, impactful research by studying both recent and foundational works, emphasizing quality over speed. Additionally, I connect them with senior researchers, broadening their exposure and fostering professional growth. This holistic approach has helped my mentees develop not only as researchers but also as thoughtful contributors to the AI community.

Undergraduate Teaching Experience. During my undergraduate studies at IIIT-Hyderabad, I served as a teaching assistant for two courses: an introductory course on “Digital Logic and Processors (DLP)” and “Introduction to Humanities.” My responsibilities included grading assignments and exams, and holding office hours to assist students. Additionally, I founded and led an Entrepreneurship Cell (E-Cell) with a social-good mission for a year. I defined its vision, organized guest lectures by socially-focused entrepreneurs, and arranged field trips to nearby villages to raise awareness of grassroots issues. A particularly memorable experience was visiting an orphanage and engaging with the children, which remains a highlight of my undergraduate journey.

Academic Service and Mentorship. I also view service activities, such as area chairing and reviewing, as teaching opportunities – offering constructive feedback and suggestions to enhance a submission rather than simply critiquing it. My dedication to this approach has been recognized with best reviewer awards at ICLR 2021, ICCV 2019, and CVPR 2017.

Future Courses

I am interested in teaching foundational courses such as “Introduction to Machine Learning” (drawing from Bishop’s and Murphy’s textbooks) and “Deep Learning” (based on the Goodfellow, Bengio, and Courville textbook), with an introduction to large language models (LLMs) and Transformer-based sequence models. Additionally, I aim to teach “Computer Vision”, covering topics from classical multi-view reconstruction to modern deep learning approaches, and an advanced course on “Vision and Language,” exploring the exciting developments over the past decade.

I would also like to offer seminar-style courses on “Multimodal AI and Foundation Models”, focusing on application-level advances, and a foundational seminar on the “Science of Deep Learning”, addressing both empirical and theoretical insights into deep learning and generalization. The latter course could include interactive demos to help students build intuition, while the former would feature a project component for hands-on prototyping.

As needed, I am also prepared to teach introductory courses such as “Statistical Learning” (based on Hastie and Tibshirani) and advanced courses like “Machine Learning Theory” (covering Shalev-Schwartz and Ben-David) and “Information Theory and Machine Learning” (drawing from MacKay’s textbook).