TEACHING STATEMENT

Pratik Chaudhari

✉ pratikac@ucla.edu   🌍 pratikac.info

Many of us can trace the decision to become engineers or scientists to an inspiring teacher. It is my ambition to be that teacher.

Courses: I can teach courses at the undergraduate level from basic probability and statistics to introductory machine learning (ML), deep learning, artificial intelligence, computer vision, robotics and control theory. I would like to develop a class at the junior undergraduate level on *"mathematics of machine learning, robotics and vision"* that introduces topics such as information theory, statistics and control while giving a flavor of applications. These are essential building blocks of modern ML but are not usually a part of the toolkit of a typical Computer Science undergraduate.

At the graduate level, I am interested in developing advanced curricula for optimization, with a focus on machine learning and deep learning, and developing the preparatory graduate course to feed them.

Experience: I was a teaching assistant (TA) for a graduate course on control theory at MIT taught by Emilio Frazzoli. I helped develop course material, hardware-based lab sessions, recitation sessions and a novel "continuous-evaluation" initiative through lightly-graded *pre*-lecture quizzes. I also have experience in teaching tutorials in the university, industry, as well as major conferences.

Education is changing: The classroom is now complemented by high-quality online material, easy-to-use code, discussions on Piazza, etc. I embrace these changes and intend to do more. For instance, a class on robotics can be designed around real robots, say a $100 quadrotor. In my experience as a TA, this gives students a taste of the real world which helps absorb the curriculum better. It also engages them on an emotional level by giving them a sense of accomplishment. I would initially leverage upon connections to the technology industry to obtain computational and hardware resources for this.

Students with very **diverse backgrounds and goals** often take courses like ML. It thus becomes challenging to develop a curriculum that can be absorbed by everyone and still remains sufficiently exciting. In my experience, two ideas are crucial to tackle this. First, continuous feedback, e.g., collecting statistics about the time spent on and the difficulty of each homework, spending time in the classroom *after* the lecture, in addition to conventional methods such as office-hours, helps to rapidly adapt the class through the semester. Second is **developing fundamentals** from the ground up instead of providing mere know-how. This works on two counts, it brings everyone on a level field and also better prepares students to explore further topics.

Research advising: Even as a graduate student, I have closely supervised junior graduate students and mentored undergraduate students. Valerio Varricchio, now a PhD student at MIT, worked with me in 2013 on autonomous vehicular fleets, which resulted in a publication at a major robotics conference. Chris Finlay, a PhD student in Mathematics at McGill University, is currently working with me on distributed optimization.

Excitement around deep learning, e.g., short-fuse ArXiv submissions and online non-peer-reviewed dissemination of ideas, can create an impression of false progress. Phenomenal results on benchmarks can be obtained with little understanding of the underlying mechanisms. Such explorations, although crucial to such a new field, do not serve the academic goal of knowledge advancement. As faculty, my role would be to identify longer-term research directions that focus on fundamentals, systematize and sharpen the thought process of my students, and ultimately, to **cultivate independent thinkers.**