# Teaching Statement

Pratik Chaudhari, University of Pennsylvania

Email: pratikac@seas.upenn.edu

I will first describe two courses that I have developed and then discuss how my teaching and mentoring methodology have evolved over the past three years. I will also discuss my outreach efforts at the end.

### New Curricula

**ESE 546: Principles of Deep Learning** I developed a new course titled "ESE 546: Principles of Deep Learning". This course begins from neural architectures, progresses through non-convex optimization, generalization of deep networks and the energy landscape of deep networks, and ends at variational inference and generative models. More than 400 students (~15% undergraduates, 70% Master's, 15% doctoral students) from across the schools of Engineering, Arts and Sciences, and Wharton have taken this course across 4 offerings. This course has received good evaluations (3.25 overall quality of the instructor and 3.25 overall quality of the course in Fall 22). Many called it the best course they ever took.

I am working with the Master of Computer and Information Technology (MCIT) program to create an online version of ESE 546 that will launch in Fall 23. It is expected to run three times a year and will reach students from across the world.

Deep learning is a very new field and lacks a good textbook. To fill this gap, I have prepared detailed notes (~200 pages, https://pratikac.github.io/pub/22_ese546.pdf). These notes have been widely appreciated by both students and faculty I have also given guest lectures on this material at Brigham Young University and Johns Hopkins University.

**ESE 650: Learning in Robotics** I took over this course in Spring 20 and revamped ~75% of the content. This is an advanced robotics course and covers state-estimation, simultaneous localization and mapping (SLAM), LQR control, dynamic programming and reinforcement learning. More than 175 students (~10% undergraduates, 70% Master's students, 20% doctoral students in Engineering) took it across Spring 20-21. Georgios Georgakis (postdoc in GRASP) and Oleh Rybkin (CIS PhD candidate) taught it in Spring 22 using my course material. This course has also received good evaluations (3.11 overall quality of the instructor and 2.87 overall quality of the course in Spring 21).

My goal is to give the students a taste of how different fields (control theory, computer vision, and machine learning) come together to create modern robotics, when ideas from these fields lead to good real-world performance, and when they fall short. I have also prepared detailed notes for this course that are well-appreciated by the students (~200 pages, https://pratikac.github.io/pub/21_ese650.pdf). I have given invited lectures on this material in MEAM 520, ESE 615, ESE 619, and also tutorials in companies (Amazon, Vanguard).

I am developing the reinforcement learning part of a popular book named "Dive into Deep Learning" (https://d2l.ai). This textbook has been adopted by 400+ universities across 60 countries and it emphasizes hands-on coding exercises while providing a broad introduction to deep learning.

### Teaching Methodology

In Fall 19, I was fortunate to be guided by prof. Santosh Venkatesh (ESE) who helped me craft my pedagogical style. I also worked with Dr. Bruce Lenthal, who directs the Center for Teaching and Learning and later took part in a semester-long Structured, Active, In-class Learning (SAIL) Seminar in Spring 22.

My goals as an Educator are as follows.

**Building foundational knowledge** Inter-disciplinary fields such as deep learning and robotics require the students to be technically broad, which may not always be the case. To mitigate this, I develop course material from the ground-up with introductory lectures at the beginning of each module (e.g., linear regression, convex functions, entropy etc.). These lectures act as the launchpad for the students and I can then focus on building a precise understanding of foundational concepts. Both ESE 546 and ESE 650 are mathematically sophisticated courses, e.g., I introduce variational calculus, steady-state distribution of continuous-time Markov chains, state-of-the-art research on understanding generalization of deep networks in ESE 546 and concepts like the HJB equation, control-estimation duality, function approximation in Q-Learning, etc. in ESE 650.

**Exposing students to research by inspiring curiosity** In both courses, students do a course project where they are free to explore a topic of their choice. I treat this as a mini-Independent Study and students receive help from their peers, TAs and me. While lecturing, I point to key open questions and help students craft course projects that investigate these questions. For ESE 650, I have procured hardware (Intel RealSense cameras and table-top robot manipulators) for students to do their course projects. Due to extensive interaction with the instructors and staff, many course projects are quite sophisticated and creative (e.g., a generative model for gene mutations; understanding the manifold of input images that is induced by a trained model; visual-inertial odometry and planning stack on an RC car). Almost all the Master's theses that I have supervised came from students I interacted with through their course projects.

**Continuous feedback** I collect statistics about the time spent on, and the perceived difficulty of each homework, inquire the students after each lecture and office hour, and run mid-semester course evaluations. This helps me adapt the course rapidly during the semester. At the end of the semester, I perform an analysis of student performance (e.g., which homework problems were difficult for students in the lower percentile, correlation of exam performance with specific homework problems etc.). Some of this analysis has proved useful to adapt the material (e.g., in Fall 19 few students grasped the content on MCMC, I have simplified it since).

**Empowering students with modern technology** I introduce cloud computing technology and provide credits to each student. In total, I have secured ~$100,000 worth credits from AWS Educate over three years. For ESE 650, I also provide tutorials on Robot Operating System (ROS) which is a standard software framework used in the industry. I have designed recitation sessions to go beyond the syllabus and provide a hands-on introduction to topics such as object detection and training large neural networks. These topics are heavily used in the industry and critical to students getting jobs/internships and being successful beyond the classroom. The exams in my courses have questions that ask students to think through a design problem (e.g., a machine learning system that is updated continually as new data is collected, imputation of missing data, selecting the sensor configuration of an autonomous forklift, etc.).

**Creating an engaging learning environment** Before every lecture, I ask students to write a 1-2 paragraph summary of the previous lecture. This keeps all the students engaged with the material throughout the semester. I often invite students to attempt to answer the question asked by someone else during the lecture. Through this, my goal is to make students comfortable enough to think out their reasoning aloud and also communicate it. While lecturing, I am very open about our poor understanding of deep learning, giving examples of my hunches that proved incorrect, or postulating hypotheses that should be explored.

**Mentoring Philosophy**

I currently advise 11 doctoral students (3 ESE, 4 CIS, 3 AMCS, and 1 BE; two are women) and 5 Master's students. Six of these doctoral students are co-advised: with Deep Jariwala (ESE), Vijay Kumar (MEAM), Christos Davatzikos (Radiology) and Jim Gee (Radiology). Yansong Gao will be the first student to graduate from my group, in Summer 23. I have supervised 6 Master's theses, served on 17 doctoral committees (9 ESE, 3 CIS, 2 MEAM, 2 AMCS, 1 Neuroscience) and 15 qualifying exam committees (CIS); I have mentored many of these students closely and have also co-authored publications with some. I expect the number of ESE doctoral students in my group to grow in the coming years.

Students I have mentored have gone on to prestigious jobs and internships in companies like Google, Amazon, Nvidia, Skydio, Nuro, Facebook, Zoox, etc., government agencies (ARL), to graduate programs/postdoc positions at MIT, Oxford, Brown, Cambridge, MILA Montreal, Purdue and Penn, and have received prestigious fellowships (NSF GRFP).

As a Mentor, I focus on the following specific goals.

**Student-centered mentoring** My emphasis is on developing good researchers with the research project acting more as the scaffolding. To do so, I work closely with students (at least 1 hr/week with each doctoral student), discuss and derive things after having them "teach" me a paper/topic (I emphasize reading books over papers at the beginning), devise experiments to cross-check their intuition, etc. I seek to develop researchers with exceptional technical depth, the wisdom to know how and why a field has evolved in a certain way, and the vision to imagine how it will evolve in the future. Many of my doctoral/Master's students had never done research before and yet they have successfully published now in premier conferences/journals.

**Promoting independence** My second goal is to prepare students to chart their own agenda and become leading researchers/academics after they graduate. After they have their first result and thereby some confidence in their own ability, I encourage them to craft the second project on their own, with feedback from me. My goal is to have them (a) develop a hypothesis based on their past theoretical and empirical experience; (b) craft the right experiment or calculation to investigate it; (c) summarize often inconclusive findings; before (d) reformulating the hypothesis. Different students take different amounts of time to appreciate this cycle and successfully go through it. I emphasize a "write-first mentality" and encourage diligent documentation as they learn to do research systematically like this.

**Creating an honest and vibrant intellectual environment** I believe that a researcher should be able to address questions of varying levels of difficulty and naivete, communicate their work, and at the end of the day inspire others to think about it. My group meetings are usually very talkative and driven by students asking questions (even early undergraduates). Addressing questions requires my students to, e.g., explain the basic idea of statistical learning theory or information geometry to a first year undergraduate, or for a computer scientist to explain to a mathematician why data augmentation is usually done on the CPU. All ongoing research progress is public within the group; every student can look at the ongoing experiments/calculations of other students. This has been extremely fruitful and many powerful ideas have come from students getting interested in each other's projects.

**Mentoring undergraduates** In the past three years, I have closely worked with more than 50 undergraduates through on-campus research programs (1st/2nd year students from PURM, VIPER, CURF, Wharton Summer Research and Google Explore Undergraduate Research, First-year exposure to STEM), Senior Design and Independent Study projects. Most of these students were conducting research for the first time. Projects ranged from hardware-based ones (SLAM using event-based cameras, active perception using a drone, oil spill containment using multiple boats, cleaning up litter using a drone) to software-based ones (estimating breast density using MRI data, generating novel Jazz solos) and theoretical ones (geometry of the predictions of deep networks, relationships between learning tasks). These projects have won multiple awards. Multiple publications have resulted from this[1–4] with one more in the works.

My primary goal while mentoring undergraduates is to give them a taste of research, put them into the right mindset for future coursework/research and give them the confidence to explore their curiosity. To achieve this goal, I meet undergraduate students on a 1:1 basis, have them take part and present in group meetings, write reports and document their experiments. Many students continue to attend my group meetings beyond the summer.

I have adapted the course material from ESE 546 and ESE 650 to initiate undergraduate students of different backgrounds to these topics. Together with the ESE Chair George Pappas and Detkin Laboratories Director Sid Deliwala, I am planning to float a 300 or 400-level course titled "Deep Learning Garage" to give students a more hands-on experience of deep learning and make them ready for more advanced coursework.

### Diversification Strategy and Outreach

To improve the diversity of my research group, I will actively recruit women doctoral students. I am a part of the ESE PhD Admissions Committee (2019, 2021, 2022) and GRASP PhD Admissions Committee (2021), which will help in doing so. Each year, I also host recruitment visits for Meyerhoff Scholars from UMBC[5] and have taken part in initiatives like Mementor[6] and Vietnam Education Foundation[7] which connect grad-school applicants from across the world with faculty in the US.

I am a co-PI on an NSF Research Traineeship (NRT) proposal across ESE/CIS, Physics, Neuroscience and Psychology. Through this, and through my existing NSF grants (CAREER and a CISE Medium grant), I hope to leverage the Greater Philadelphia Louis Stokes Alliance for Minority Participation (AMP) to recruit students.

In Fall 23, I am planning to organize a Greater Philadelphia Machine Learning Day to create a local community of researchers in academia and industry who work on machine learning. This workshop will be co-located with a hackathon that is run by PennApps each year (which is the nation's largest student-run hackathon). We hope to interest over 1000 students (there were 1251 participants in Fall 19 and 974 in Fall 20 at PennApps from across the US). My goal is to have a different local institution host this event each year (similar to the New-England Robotics Colloquium that I co-organized in 2019); multiple institutions have expressed interest in doing so. As a part of local outreach efforts funded by my NSF CAREER grant, I am planning to work with the Franklin Institute to organize after-school-hours hands-on tutorials for students in grades 8–12 (Python programming, face recognition using the laptop camera, sentence completion etc.), and through this explain how physics and math play a role in building AI systems.

### References


1. Lakshay Sharma, Aditya Hota, Zhifei Shen and Pranav Panganamamula. GRASPer: A Drone-Based Automatic Litter Collection System. in *4th Annual Andrew P. Sage Memorial Capstone Design Competition (SMDC)* (2021).
2. Garrett, A., Duhamel, J., Liang, A., Goldberg, Z. & Friedman, J. BoomBoat: A Multi-Robot Automated Oil Spill Containment Solution. in *5th Annual Andrew P. Sage Memorial Capstone Design Competition (SMDC)* (2022).
3. Cowlagi, A. & Chaudhari, P. Does the Geometry of the Data Control the Geometry of Neural Predictions? in *AAAI (Student Track)* **36** (2022), 12931–12932.

4. Katti, K., Muthukrishnan, R., Heyler, A., Pati, S., Alahari, A., Sanborn, M., Conant, E. F., Scott, C., Winham, S., Vachon, C., Chaudhari, Pratik, Kontos, D. & Bakas, S. MammoDL: Mammographic Breast Density Estimation Using Federated Learning. *arXiv preprint arXiv:2206.05575* (2022).
5. Meyerhoff Scholars Program. https://meyerhoff.umbc.edu/.
6. Mementor Beta. https://mementor.net/#/.
7. VEF 2.0 Program |. https://www.vef2.org/.