I love helping students have their 'aha!' moment, either when learning a new fundamental concept through a course or when solving problems for their research projects. I strive to create an environment where junior folks feel intellectually safe to ask all of their questions, without fear of sounding stupid, and keep drilling down on problems until they fully figure them out. My teaching philosophy is built on three core principles. First, I focus on creating a **deep, expandable understanding** that goes beyond single instances of concepts, helping students form mental maps that allow them to apply ideas across diverse contexts. Second, I foster an **intellectually safe space** that encourages questions and disagreements without judgment, promoting curiosity and critical thinking. Third, I strive to make education **accessible** by adapting teaching styles to suit different learners and mentees, and by providing supplementary materials and resources that are made publicly available. I strongly oppose having a single, fixed teaching/advising style since I believe there is no one-size-fits-all approach to education. I maintain that everyone is capable of learning, and it is on us, the educators, to identify and apply the best advising or teaching style for each person to make the information accessible to everyone.

These principles have been shaped and reinforced through my diverse teaching experiences, from high school volunteer work to graduate-level courses and mentorship programs. In the following sections, I elaborate on these experiences, demonstrating how they have influenced my approach to education and prepared me for future teaching roles.

1 Teaching Experience and Philosophy

Throughout my graduate studies at UC San Diego, I actively sought out teaching opportunities to refine my skills and make a positive impact on students' learning experiences. I served as a teaching assistant for six graduate and undergraduate courses, including Intro to Machine Learning, Mathematics for Robotics, and Computer Architecture. I have also been asked on numerous occasions to deliver tutorials, guest lectures, and invited talks on my research since I am one of a scant number of domain experts on privacy, safety, and large language models:

- Guest lecture for CMU LTI 11-830 course on Computational Ethics in NLP
- Guest lecture for UW CSE 484 and 582 courses one Computer Security and Ethics in AI
- Guest lecture for UT Austin LIN 393 course on Social Applications and Impact of NLP
- Guest lecture for Johns Hopkins CS 601.670 course on Artificial Agents
- Tutorial co-instructor for EACL 2023 tutorial on Private NLP
- Keynote speaker for EthiCS workshop at NDSS 2023
- Invited presenter at Generative AI and Law workshop at Washington DC on differential privacy

One of my primary goals as an instructor is to help students create connections between different topics, building a comprehensive knowledge graph rather than storing isolated facts. For instance, when discussing membership inference attacks in the context of privacy, I challenge students to think beyond the immediate application: I ask them to consider other ways they can use this tool, and then I walk them through how this can be adapted for attribution and copyright violation detection or even test set contamination, which are seemingly unrelated issues. My teaching experience is not limited to my graduate studies. As an undergraduate, I volunteered as a TA for eight courses, including Digital Electronics, Computer Architecture, Signals and Systems, Probability and Statistics, and Numerical Methods. I taught some of these courses multiple times, which helped me to refine my teaching methods and develop a deeper understanding of how to effectively communicate complex concepts.

2 Mentoring Experience and Philosophy

My approach to mentoring is guided by the principle of 'teaching how to fish' rather than simply providing answers. I strive to equip students with the skills and mindset necessary to navigate the challenges of research and academic life independently. This includes fostering creativity, teaching how

to build narratives and tell compelling stories about their work, encouraging students to follow their interests rather than simply chasing trends, and navigating possible discomforts and uncertainties that arise when doing research.

I began formal mentoring in 2020 through OpenMined's volunteer research program, which connects privacy enthusiasts from around the world. This experience allowed me to mentor four students on projects focused on differential privacy, fairness, and interoperability, resulting in publications at major conferences. Since then, I have expanded my mentoring efforts, advising 22 students on long-term research projects from various institutions both within and outside the US. Five of the undergraduates I mentored have gone on to pursue PhDs at prestigious institutions such as Carnegie Mellon, Georgia Tech, and MIT. The impact of my mentoring is reflected in the words of one of my mentees:

> "Hi Niloofar, I received my first PhD admit from Georgia Tech to work with Prof. [redacted] and Prof. [redacted]!! Thank you so much for everything, I'm totally indebted to you, for all the feedback/mentorship. You inspire me!"

Beyond research guidance, I have also advised students on PhD applications and often-overlooked aspects of academic life, such as mental health and work-life balance. In 2021, I initiated a Twitter-based program pairing students with mentors for feedback on statements of purpose and PhD applications. I have also volunteered in various mentorship programs, including roles at ICLR 2021, the Women in Machine Learning workshop at NeurIPS 2020, and the UC San Diego Women Organization for Research Mentoring in STEM.

3 Course Design and Future Plans

Looking ahead, I am enthusiastic about developing and teaching a diverse range of courses in artificial intelligence (AI), machine learning (ML), natural language processing (NLP), and societal implications of AI. Given the rising importance of safety and governance for large language models (LLMs), I plan to design a curriculum focused on LLM privacy and security. This course will build upon the module I created for the UW CSE 447 NLP course.[1] The structure will blend instructor-led sessions with student-driven research paper discussions and idea generation. In paper discussion sessions, students will prepare proposals for experiments that expand on or test additional hypotheses related to the paper's content. To foster innovative thinking during these discussions and brainstorming sessions, I intend to employ creative methods such as Oblique Strategies.[2]

My approach to homework assignments is designed to deepen students' understanding through hands-on learning. Assignments will guide students through step-by-step implementations of algorithms covered in class, allowing them to construct and analyze different components. This method enables students to gain insights into how various elements interact since I believe that true, sustainable learning occurs through active problem-solving. For example, when teaching fundamentals of privacy in machine learning, I would begin with a practical scenario: identifying smoking habits in a community. After soliciting student proposals, I would introduce the randomized response algorithm. Subsequent homework assignments would challenge students to generalize this algorithm beyond binary cases and eventually apply differential privacy concepts to simple ML models. These assignments would then form the basis for upcoming class discussions.

In my instruction, I strive to provide a contextual narrative for each concept or method. This includes explaining the historical context of its development, the rationale behind chosen datasets and evaluation metrics, and how these choices have influenced subsequent research. For instance, when discussing memorization, I would highlight the different evaluation approaches used in privacy papers versus interpretability studies. My goal is to equip students with the analytical tools to disentangle these nuances and construct deeper conceptual frameworks. By offering this comprehensive and contextualized approach to learning, I aim to cultivate a rich understanding of both the theoretical foundations and practical applications of computer science concepts among my students.

[1] https://courses.cs.washington.edu/courses/cse447/24wi/

[2] https://en.wikipedia.org/wiki/Oblique_Strategies