Teaching prepares students for a career in research while mentorship is crucial for their success. I have had the privilege of having excellent teachers and supportive mentors throughout my career. They have shaped my philosophy in these crucial aspects of being an academic, allowing me to be both rigorous and empathetic. Both teaching and mentorship give me a great sense of fulfillment and have led to me pursue being a professor as my lifelong goal.

Teaching

Experience. My first teaching experiences were at the Indian Institute of Technology Madras, where I served as Teaching Assistant for a graduate-level theoretical course on Convex Optimization and a lab on Communications. My responsibilities in the former largely consisted of grading assignments. The lab course was hands-on, and I had to help students with experiments. I enjoyed the practical aspect of teaching this course, as helping students to ensure their experiments worked was very rewarding and enhanced my hands-on skills. At Princeton University, I was a teaching assistant for a challenging graduate-level course on the theoretical underpinnings of machine learning (ML) that was also open to undergraduate students. In spite of the rigor of the course, it attracted students from diverse backgrounds as interest in ML was burgeoning. During office hours, I often ended up re-teaching concepts covered in class which I found rewarding as it deepened my understanding of the source material. It also allowed me to expose motivated students to open research questions, as these often arise when probing the fundamentals of any topic. I also learned to guide students to the correct answer without giving away the solution entirely, to develop critical thinking and theorem-proving skills.

I recently held a 2-day workshop introducing students from diverse backgrounds and skill levels to the fundamentals of data science. I used an interactive Python notebook to conduct the sessions. I created a teacher and student copy of the notebook, which in itself was an interesting exercise in determining what concepts the students had to grasp on their own, and which could simply be shown. I had to juggle live questions from students with interests ranging from the purely practical to the more theoretical, while ensuring everyone was engaged. The examples used in the teaching material concerned the analysis of child education indicators across countries, which was an ideal case study to demonstrate both the benefits and limits of quantitative approaches for solving problems of social relevance.

Philosophy. I believe that teaching is an act of giving, with the teacher imparting knowledge with *rigor* and *empathy*. This must be done while instilling a sense of *curiosity*, while maintaining *humility* with an acknowledgment of the limits of the teacher's knowledge.

Rigor and Curiosity. When teaching, I aim to be rigorous, instilling a deep understanding of fundamental concepts while relating them to practice and application. This can be done with the effective use of technical aids to teaching, with simulations and coding exercises complementing pen-and-paper assignments. A rigorous understanding can be cultivated with regular assignments that challenge the students, in lieu of a heavy reliance on tests. I will pique the students' curiosity by demonstrating how even the most abstract of technical concepts find application, thus appealing to both the theoretically- and practically-minded students. In class, I will conduct interactive lectures that encourage students to answer questions and make logical leaps. Certain conclusions can be left unsaid for the students to discover in an assignment or in discussion groups.

Empathy. Cultivating curiosity goes hand-in-hand with an empathetic approach that makes students feel comfortable asking questions, even if they are unsure of their question. The human aspect of teaching is as important as the technical, so I will take an empathetic approach in the classroom. I plan to solicit feedback from students on a regular basis and adapt the course material and delivery if necessary. I will ensure that all students feel welcome in my class, regardless of caste, gender or religion to foster an optimal learning experience. In practical terms, this will involve making the course expectations clear in the beginning of the semester, with a clear syllabus, grading rubric and policies around the use of aids such as generative AI tools. I will allow a fixed number of late days for assignments, as well as the possibility for make-up exams, depending on the size of the class. For courses taught to students in their third year or later, I will focus on designing open-book examinations as I have found that these test conceptual understanding much better than closed-book exams that tend to reward memorization.

Intellectual Humility. Finally, I seek to practice intellectual humility and inculcate it in my students as well. I will position myself as a guide to knowledge open to being questioned, and encourage students to go beyond the material covered in the class. In practice, I will hold regular office hours and work closely with teaching assistants to ensure they are equipped to handle students' questions and concerns.

Potential courses. My teaching and research experience has prepared me to teach a range of courses in computer science and electrical engineering. I could teach undergraduate courses on data science and machine learning, com-

puter security, probability and statistics as well as signal processing. For more advanced undergraduates and graduate students, I could teach courses on deep learning, including new topics such as distributed learning and generative models as well as convex optimization. I also keen to design a new course on failure modes of machine learning (ML), exposing advanced students to the myriad ways in which ML models can fail, which is a burgeoning new area of research with many open questions. I am prepared to design such a course based on my contributions to two book chapters [1, 2] and the key monograph on federated learning [8]. Working at IBM Research gave me a perspective into ML research in industry and I look forward to organizing guest lectures by people working on ML at different industries to help students understand how ML is applied in practice.

Mentorship

Experience. During the course of my Ph.D. and post-doctoral experience, I have mentored 12 Ph.D. and Masters, 3 undergraduate and 2 high school students. I am glad to have been able to mentor students from diverse backgrounds, across nationalities, age and gender. These mentorship stints have led to several co-authored papers with these students at top-tier conferences in computer science [5, 4, 11, 13, 12, 10, 9, 3, 7, 6]. For me, working with students is the most exciting part of research, as I thrive in a collaborative research environment. One of my greatest joys is being able to transfer my excitement about a line of research or new idea to a passionate student and watch them develop their own mode of thinking about the problem. I have also served as an indirect mentor, advising Ph.D. students on how to cultivate their own mentoring relationships.

Philosophy. My mentoring philosophy rests on three key principles: *respect*, *rigor* and *responsibility*. Together, I believe these lead to a collaborative and congenial atmosphere where challenging problems that benefit society can be tackled.

Respect. The cornerstone of healthy professional relationships is respect. Research is an intellectual endeavor that benefits from points of view, which can often lead to interpersonal friction. However, any critiques must be expressed constructively, with the aim of benefiting the project. For good research to happen, there must be a sense of ownership over the work and respect for the problem being tackled. I seek to cultivate this approach in all the students I work with, so they are respectful of both the work as well as the individuals undertaking it.

Rigor. Research in computer science benefits greatly from the application of fundamental scientific principles. I encourage my collaborators, particularly students I work with, to take a clear-minded approach to setting up a well-motivated problem, with a crisp understanding of where the work is located with respect to existing literature. I further urge them to continually query both the motivation and method used to tackle the current problem, which can lead to creative new solutions and approaches. Honesty in carrying out research as well as in reporting is key to sustaining good scientific work.

Responsibility. As computer scientists,it is our prerogative to undertake meaningful research that minimizes harm. I urge the students I mentor to think about the implications of their work, particularly as it focuses on AI and ML, which are technologies that are getting embedded in the fabric of our day-to-day lives. Responsible citizen scientists can help mitigate any ill effects from their research in the future. I aim to cultivate a sense of responsibility towards the next generation of researchers among my mentees in order to ensure a virtuous cycle of mentorship continues.

Practice. In order to implement the goals of my mentorship philosophy, I will work closely with my students in both an individual and group setting. I will hold regular weekly meetings for one hour with each student in order to track their progress and address any concerns they may have. I will strive to make this an expectation-free, safe space where the students and I can productively discuss their research, and on occasion, other concerns they may have. Having the confidence to express their ideas is the first step for students to become good researchers. I will first encourage students to propose their own projects, which I will then review in detail with them to ensure the motivations and method are sound. I will also keep a running list of potential projects for students to choose from as an alternative. I will also hold a group meeting for ninety minutes each week which all students and postdocs will be expected to attend. Each meeting will have a clear agenda, whether discussing ongoing work in the group, doing dry runs for external presentations or doing deep dives into relevant literature. I will also support and encourage student-driven reading groups that are needed to stay current with ongoing research. Research in ML moves very quickly, and it important for students to be aware of the latest developments. I will encourage my students to maintain a healthy work-life balance and organize occasional team-building events for the group. I plan to maintain my current collaborations with academics across the United States as well as in industry, and will actively seek to include my students in these to provide exposure to world-class researchers. Finally, I will train and urge the group to use communication and productivity tools like Zotero, Slack, and Github.

References

[1] **A. Bhagoji** and S. Chakraborty. Securing federated learning: Defending against poisoning and evasion attacks. In L. M. Nguyen, T. N. Hoang, and P.-Y. Chen, editors, *Federated Learning: Theory and Practice*. Elsevier, 2024.

[2] **A. Bhagoji** and P. Shirani. Adversarial attacks on anomaly detection. In D. Phung, G. I. Webb, and C. Sammut, editors, *Encyclopedia of Machine Learning and Data Science*. Springer US, 2020.

[3] J. Brown, X. Jiang, V. Tran, **A. Bhagoji**, N. P. Hoang, N. Feamster, P. Mittal, and V. Yegneswaran. Augmenting rule-based dns censorship detection at scale with machine learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2023.

[4] C. Cianfarani, **A. Bhagoji**, V. Sehwag, B. Zhao, H. Zheng, and P. Mittal. Understanding robust learning through the lens of representation similarities. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS)*, 2022.

[5] S. Dai, W. Ding, **A. Bhagoji**, D. Cullina, B. Y. Zhao, H. Zheng, and P. Mittal. Characterizing the optimal 0-1 loss for multi-class classification with a test-time attacker. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (NeurIPS)*, 2023.

[6] W. Ding, **A. Bhagoji**, B. Y. Zhao, and H. Zheng. Towards scalable and robust model versioning. In *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 2024.

[7] X. Jiang, S. Liu, A. Gember-Jacobson, **A. Bhagoji**, P. Schmitt, F. Bronzino, and N. Feamster. Netdiffusion: Network data augmentation through protocol-constrained traffic generation. *Accepted with Shepherding in Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2024.

[8] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, **A. Bhagoji**, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 2021.

[9] S. Liu, F. Bronzino, P. Schmitt, **A. Bhagoji**, N. Feamster, H. G. Crespo, T. Coyle, and B. Ward. Leaf: Navigating concept drift in cellular networks. *Proceedings of the ACM on Networking*, 2023.

[10] A. Panda, S. Mahloujifar, **A. Bhagoji**, S. Chakraborty, and P. Mittal. Sparsefed: Mitigating model poisoning attacks in federated learning with sparsification. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.

[11] V. Sehwag, **A. Bhagoji**, L. Song, C. Sitawarin, D. Cullina, M. Chiang, and P. Mittal. Analyzing the robustness of open-world machine learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security (AISec)*, 2019.

[12] E. Wenger, R. Bhattacharjee, **A. Bhagoji**, J. Passananti, E. Andere, H. Zheng, and B. Zhao. Finding naturally occurring physical backdoors in image datasets. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS)*, 2022.

[13] E. Wenger, J. Passananti, **A. Bhagoji**, Y. Yao, H. Zheng, and B. Y. Zhao. Backdoor attacks against deep learning systems in the physical world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.