Teaching Statement

Shivaram Venkataraman

The opportunity to teach students and collaborate with them on research projects is one of the primary reasons I aim to work in academia. I have greatly enjoyed teaching and designing course materials while in graduate school and I look forward to continue contributing to computer science education in the future.

Classroom Teaching

Experience: I have had teaching experience during my time in college, masters and PhD. Through my experience, I have had a chance to explore various aspects of designing, organizing and teaching courses.

While I was an undergraduate student at the Birla Institute of Technology and Science, Pilani I was a lab assistant for the introductory computer programming course and helped sophomores with their lab assignments. As a masters student at UIUC, I was a teaching assistant across two semesters for the undergraduate course on data structures typically taken by sophomores and juniors. In addition to teaching sections I also helped create exam problems, and being a part of the teaching staff for a class with around 200 students helped me learn how to organize a large course.

At UC Berkeley I was a teaching assistant for the database course typically taken by juniors and seniors. This was similarly a large class with around 300 students. As a part of teaching this course I also proposed updating one of the projects to cover more advanced database operators. To do this, I designed a new project that introduced probabilistic data structures and how they can be implemented within a real-world database like PostgreSQL. Along with my advisor Mike Franklin, I was also involved in designing a new course introducing data science to CS students. Starting from a clean slate allowed us to select both *what* material was taught and also *how* the material was taught. As the data science curriculum spans across statistics and computing, we interleaved lectures from both areas. In order to train students on the latest tools used for data science, I helped create hands-on exercises that were part of every lecture.

Finally, for the past two years I have presented guest lectures in the advanced parallel computing course taught by Prof. Jim Demmel. These lectures introduced cloud computing and big data processing to graduate students from various departments with different backgrounds.

Courses I Can Teach: Based on my experience, I can teach undergraduate and graduate level courses in **databases**, **operating systems**. I am also interested in teaching an undergraduate **applied machine learning** course on how practitioners can effectively use machine learning. Finally, I would like to teach a graduate-level seminar on recent advances in my research area covering big data analytics, cloud computing and applied machine learning.

Additionally given the emerging importance of **data science** I would like to contribute to developing two courses in this area. The first course would be targeted at CS majors, especially juniors or seniors who are proficient in CS topics and would like to learn how to apply statistical concepts in data analysis. This course would introduce applied machine learning topics like classification, feature selection and clustering. The course will also include database topics like data cleaning, indexing and how a mixture of classical data management techniques can be combined with statistical techniques. The second course would be targeted at CS freshmen or more broadly across majors and would cover the foundations of data science. This will include basic statistical concepts like sampling and regression, and also incorporate computational topics like numerical libraries and visualization tools. For both courses I hope to incorporate projects that use real-world data analysis scenarios as I have found that such hands-on experience is very useful for understanding various aspects of statistical modeling. To do this I would like to explore opportunities for students to work with data from non-profits organizations, local government agencies or other campus units.

Teaching Approach

The growth of jobs in data science and software engineering and the rapid changes in technology pose new challenges to CS education. In particular, we need to design courses such that students can learn how to adapt to changes in technology and handle growing class sizes.

Through my teaching experience in graduate school I have found that the most important role of computer science education is to foster understanding of the core principles of computing. Having a sound understanding of foundational principles, coupled with a understanding of how those principles are reflected in real-world scenarios enables students to pick up new technologies through their careers. In particular I believe this can be only be achieved by having a mix of theoretical concepts and hands-on practical exercises. For example, in the data science course at Berkeley, we found that the lab exercises were useful not just for improving coding skills, but were also crucial to improve understanding of statistical concepts by say visualizing data or measuring the utility of feature selection. Thus the students got a chance to see how the theory they learned was reflected in practical scenarios.

With the growing popularity of CS courses, I also believe it is important to enrich the course structure to make sure that they handle students with a wide range of abilities. For example, while teaching the data structures course at UIUC, I designed extra projects that more advanced students could take up as a part of an honors course. Further, to create an atmosphere where students are more comfortable asking questions, along with Prof. Cinda Heeren, I organized a new program called 'Own the Final'. In this program, we created small study groups that students could voluntarily join. We encouraged participation by asking specific students to prepare review material for each session and found that students benefited a lot by learning from each other. Overall this program was successful in helping students prepare for the final exam and some students improved by up to two letter grades.

Research Mentoring

During my PhD, I have also had the opportunity to work with many talented undergraduate students and guide their research. I typically work closely with the students and thus attempt to integrate them into my current research projects. At the same time I tailor their assignments to match their goals and aim to teach them engineering and research skills that will be useful in their future.

For example, Zongheng Yang, a sophomore then a UC Berkeley was interested in pursuing research topics and I worked closely with him on the SparkR project. Zongheng made several significant contributions to the project and I encouraged him to present talks about his work at Spark Summit and AMP Camp. Zongheng continued working with me on Ernest, a research project to do efficient performance prediction. He is currently at Google and will be starting his PhD at UC Berkeley next year. In the case of another student, Tomer Kaftan, I initially helped him develop software development skills like refactoring code, writing design docs. As Tomer developed interest in contributing to research, I helped him become a part of the research on KeystoneML and he worked on the software architecture for the KeystoneML optimizer. Tomer is now a PhD student at the University of Washington.

In summary, I believe that using a hands-on approach and providing regular feedback allows students to make rapid progress on their research. With my experience in industry, I also help them learn good software engineering practices and thus provide a foundation for students pursuing industrial or academic positions in the future.