TEACHING STATEMENT

SAMUEL B. HOPKINS

Algorithms, statistics, and data science are suddenly in the public eye and at the forefront of students' minds, as evidenced by ballooning enrollments in computer science, machine learning, and statistics courses. It is an exciting time to teach any of these disciplines! I am eager to teach courses at all levels which combine algorithmic and statistical thinking, along with the traditional algorithms and complexity curriculum. As I work to develop my teaching in the coming years, I aim to be guided by a few principles:

- *Curiosity.* I aim to foster students' curiosity and intellectual excitement via discussion and opportunities for student-driven projects.
- *Context.* The material we teach, even in theoretical courses, does not exist in a vacuum. It was invented to solve real and important problems – we should teach in historical context. It has broad social and scientific impact – we should teach in social and scientific context.
- *Evidence.* I aim to incorporate evidence-based approaches to teaching into my courses.

Experience. As a PhD student at Cornell, I was a teaching assistant for two courses: compilers and computational complexity, both taken by 3rd and 4th year undergraduates. To teach compilers I had to leave my comfort zone – at many points in the course, I learned the concepts being taught just a week or two ahead of the students. Although I would never again tolerate such lack of preparation, the experience gave me the confidence that with humility and grit on the part of the instructor it is possible to teach outside your immediate area of expertise when called upon to do so.

As an undergraduate student at the University of Washington, I worked as a tutor for two years at UW's Philospohy writing center and as a TA for undergraduate probability. As a writing center tutor, I held numerous one-on-one meetings with undergraduate students to discuss their in-progress essays. These meetings taught me that that listening to students is just as important as talking to them. Listening is not just part of creating a welcoming and inclusive environment: it is also the only way to pinpoint exactly what concept a student may be struggling with.

Mentoring. As a postdoc at UC Berkeley, I have mentored a number of graduate students from across the theory of computing, machine learning, and biostatistics groups. This has led to several long-term projects in which I have acted as the primary senior member of the team – I worked with students to choose problems, suggested mathematical techniques, and taught students relevant background material. One of these projects, involving students from the theory of computing and machine learning groups at Berkeley, has resulted in a substantial paper in heavy-tailed statistics, currently in submission to STOC 2020 (one of the premier publication venues in theory of computation) [C**H**KRT19].

Proposed Courses. I am excited to teach courses in discrete math, theoretical computer science, and theoretical machine learning and data science, at the undergraduate and graduate levels.

Especially at the undergraduate level, the machine learning revolution has led to changing needs for courses in algorithms, both to ensure that students grasp some of the algorithmic underpinnings of machine learning and to ensure that they have the skills needed in industry. I propose to teach algorithms courses which are *ML-aware*, while still preserving fundamental material such as dynamic programming and network flow. To me, this means that the algorithms course should contain a beginning treatment of gradient descent and at least one example of beyond-worst-case analysis.

At the graduate level, there is ample opportunity to infuse the algorithms curriculum with basic algorithmic principles relevant to data science and machine learning. I would aim for graduate algorithms students to leave the course with some appreciation of the amazing power of convex optimization and spectral methods in both worst-case and beyond worst-case settings. This could mean teaching compressed sensing, matrix completion, sparse PCA, or algorithmic robust statistics in addition to traditional topics in graduate algorithms.

REFERENCES

[C**H**KRT19] Yeshwanth Cherapanamjeri, Samuel B Hopkins, Tarun Kathuria, Prasad Raghavendra, and Nilesh Tripuraneni. "Algorithms for heavy-tailed statistics: regression, covariance estimation, and beyond". In: *In submission* (2019).