

I genuinely enjoy interacting with students, both through teaching and research mentoring. I have a great deal of experience in both of these areas, which will allow me to hit the ground running as a new faculty member.

During the Fall 2021 semester, I was the co-instructor of 15-445/645: Database Systems [1] at CMU, an upper-level systems course about the internals of a database management system (DBMS). My responsibilities covered all aspects of course development, including: (1) creating a syllabus and course plan; (2) managing teaching assistants; (3) developing exams, programming projects, and homework assignments; and (4) preparing and presenting lectures. The course had an enrollment of 130 students, comprised evenly of undergraduate and graduate students.

As I explained in my [Diversity Statement](#), I believe that research and teaching products should be made accessible to the widest possible audience. All of my course materials for 15-445/645 are freely available on the course website [1] (including lectures [2]), and I am committed to continuing this practice for every course I teach in the future. I also plan to explore other ways of making course content more accessible, such as by adapting the 15-445/645 materials into a mini-course with condensed lectures that might be more digestible for non-students (e.g., full-time workers).

In addition to 15-445/645 at CMU, I have had experience as a member of course staff for data management courses at two other universities. As a Ph.D. student, I was the head teaching assistant for CSCI 1270: Database Management Systems [3] at Brown, which was also a DBMS internals course. Although I gave a few lectures during the semester, my primary role was to write the exams, programming projects, and homework assignments, as well as to lead the undergraduate teaching assistants in grading. In particular, I redesigned several of the existing programming projects, which had students implement one-off components (e.g., a B+tree), so that they were integrated into a full-fledged pedagogical DBMS called SimpleDB [6]. These assignments allowed students to understand how the different components fit together at the level of the entire system. I believe that the easiest way for students to grasp complex concepts in systems courses is through direct hands-on learning, so I plan to continue this model in the future courses that I teach.

As an undergraduate, I was a teaching assistant for CSCI 2257: Database Systems and Applications [4] at Boston College for two semesters. Unlike the other two DBMS internals courses, this course focused primarily on the user-facing aspects of a DBMS, including introductory SQL, application development, and database administration. These topics naturally attracted students from areas outside of computer science (e.g., business, physical sciences) who wanted to learn how to effectively use a DBMS in other domains. In fact, the entire sub-field of data management has always been strongly driven by understanding and responding to the unique needs of end users. Therefore, I believe it is essential to offer these types of courses in order to make our work more accessible, and I have experience tailoring course materials to an audience without a strong computer science background.

Beyond course development, I also have significant experience with research mentoring. Since the beginning of graduate school, I have worked closely with 36 students (5 Ph.D., 10 master's, 21 undergraduate) on a wide range of research projects. For Ph.D. students, I mainly provided feedback and helped them to develop research ideas related to their dissertation topics. On the other hand, my experience with undergraduate and master's students was much more varied; some contributed to my ongoing projects (e.g., prototyping, running experiments), whereas others worked on standalone projects. I have had 8 undergraduate and master's coauthors on full research papers published in top-tier conferences, and two papers currently under submission were led entirely by undergraduates.

As a new faculty member, I would be excited to teach a variety of data management courses at all levels, including user-facing applications, system internals, and advanced concepts. Given my research focus and involvement with the Data Science Initiative [5] at Brown, I am also well-positioned to teach introductory data science courses. Lastly, I am looking forward to an opportunity to develop a graduate-level seminar course related to my research expertise: rethinking the design of traditional data analytics systems to fully leverage emerging trends in hardware and workloads.

[1] <https://15445.courses.cs.cmu.edu/fall2021/>

[2] <https://www.youtube.com/playlist?list=PLSE8ODhjZXjZaHA6QcxDfJ0SIWBzQFKEG>

[3] <http://cs.brown.edu/courses/csci1270/>

[4] https://www.bc.edu/bc-web/schools/mcas/departments/computer-science/academics/courses.html#tab-two_thousand_level

[5] <https://www.brown.edu/initiatives/data-science/home>

[6] E. Sciore. SimpleDB: A Simple Java-Based Multiuser System for Teaching Database Internals. In *SIGCSE*, pages 561–565, 2007.