Xu Wang Teaching Statement

As an educator, I have four overarching goals towards teaching: developing courses to deliver the knowledge we have produced in research; deploying the technologies from my research to support teaching and learning; mentoring students and empowering students to begin research early in their academic careers; facilitating the development of online programs to enable more people get access to high quality educational resources.

Teaching

My teaching and research experience has informed each other. I was a teaching assistant for CMU's User Centered Research and Evaluation (05610/410) course in Fall 2017. My responsibility included *leading lectures and interactive sessions weekly to a group of 20 students, designing assignment questions and prompts, grading and offering individual feedback to students, and holding weekly office hours*. I was also a teaching assistant for CMU's Applied Machine Learning (05834) course in Spring 2018. This is a course designed for students who do not have strong programming background. Office hours became a critical venue for them to catch up on content such as probability. I was very excited to teach machine learning concepts to students who were not in the field, and my office hour was always full and sometimes extended due to the large number of students who came to seek help.

I apply effective teaching techniques known from the learning sciences literature and my own research into teaching. I use **worked examples** and facilitate **learning by doing practices** during my lectures and sessions. For example, for the class session on "Survey Design", I wrote positive and negative examples of survey questions. I asked students to evaluate each example survey question and find out the issues with them. I then went over each example with the entire class. This class session was well received because it uses concrete examples to consolidate the task of writing good survey questions. During my lecture and office hours, I try to **elicit students' prior knowledge and understand their thinking** rather than directly tell them the right answer. Moreover, **I use student data to tailor my instruction.** For example, for the class session on "Naive Bayes classifier", I found many students who did not have a solid background in probability had trouble understanding that what they observe is by a certain distribution of probability, and they need to model the chances that what they observed had happened. I thus tailored my instruction during office hours to focus on basic concepts around probability before delving deeper into Naive Bayes classifiers.

Support Teaching and Learning with Technologies

My teaching practice has informed my research on developing technologies to support teaching and learning. I noticed that large complex open-ended assignments were often given, many student submissions shared common misconceptions, and the same error often repeatedly occurred within one submission. As class size gets bigger, feedback to students is often delayed. This does not align with the learning science research that suggests that deliberate practice and immediate feedback helps with learning. To address this gap, I developed UpGrade, a system that leverages past students' answers to create multiple-choice practice questions. In contrast to traditionally used open-ended assignments, UpGrade-produced practice questions offer real-time feedback and repeated practice. We have demonstrated that UpGrade can increase student learning efficiency and reduce efforts from instructors. **UpGrade has been used in 7 modules of 8 classes at CMU, including User Centered Research and Evaluation, Programming User Interfaces, E-Learning Design Principles and Research Methods for Learning Sciences**.

With UpGrade as a starting point, I am excited about bringing technologies to support teaching and learning within the university, and building online programs to extend our educational impacts. I am interested in creating better reusable infrastructure for courses, where student data is better structured, archived and reused; instructor effort is optimized and repetitive work is reduced to minimal. Additionally, I hope to facilitate the sharing of instructional materials across courses and institutions. Building the platform to support teachers collaboratively

design and improve courseware could reduce required efforts and help craft higher quality materials over time. Moving forward, I have a strong interest in facilitating the development of online programs to allow more people get access to high quality educational resources. The introduction and development of such infrastructure could 1) support the delivery of higher quality learning materials to students; 2) reduce repetitive work from faculty members; 3) extend higher education opportunities to a larger community; 4) and enable data-driven instruction and assessment design, with potential practical and research impacts.

Mentoring

During graduate school, I have been fortunate to mentor 2 undergraduate and 4 graduate students at CMU in various research projects. Yali Chen worked on a project to improve the quality of feedback being exchanged in online learning environments. She developed a coding framework to assess the quality of feedback, which is still in use in other projects of our lab. Her work resulted in a second-author full paper at ICLS 2018. Yali went to on to become a product manager at Cengage, which is a leading company on educational content development. Teja Talluri applied and refined psychometric methods to detect unreliable question items from an auto-generated question bank. His work on quality control of UpGrade resulted in a second-author full paper at L@S 2019. He continued to investigate the use of psychometric methods to benefit small to medium size classroom instruction, with a publication under preparation. Teja is now a data scientist at Inspire, where he uses computational methods to support online health communities. Kexin Yang worked on a project in collaboration with the MIT Teaching Systems Lab on a teacher training platform. She designed and implemented both qualitative and quantitative methods to analyze the conversation data we collected from two experiments. Her work resulted in a third-author full paper that is currently under submission. Kexin got very excited about educational research and she is now a PhD student in the school of Computer Science at the University of Waterloo.

I use a combination of hands-on and hands-off techniques for mentoring. For example, early on in the research project, I tend to offer more scaffolding and provide concrete tasks for the student to explore. I usually fade my guidance as the project develops. I also constantly check in with my mentees to make sure they are getting the help they need and are making progress. Having constructive conversations between me and my students has been a successful strategy that I will continue to use in my career. Such diverse and successful mentoring experiences were highly rewarding, which motivates me to continue to help my students grow as independent researchers.

Community

Throughout my career, I have been active in introducing computer science and HCI to a broad audience, from middle school girls to researchers in other fields. During my time at CMU, I participated in several events teaching high school girls computer science, including a session in the CMU TechNights program (link) in which I taught Scratch to middle school girls, and a Community STEM Night event at the South Side Middle School in Pittsburgh in which I demoed robots and explained recommender systems to kids and parents. I also served as a mentor for the CMU OurCS program (link) in 2018, which is a three-day workshop to encourage undergraduate women to explore research. I was a mentor for a group of 20 students, supporting them to use machine learning toolkits to tackle social challenges.

I am excited about building interdisciplinary research communities. I have been an associate with CMU PIER: Program in Interdisciplinary Education Research (link), where doctoral students from a variety of schools, including computer science, psychology, business, public policy, etc., all come together and share research. It has been a perfect platform for me to get feedback from different perspectives and share my research to a broad community. I was an organizer for the PIER Edbag events for a year, for which we invite CMU student researchers and faculties to give talks or lead discussions. I also co-organized three PIER seminars, which are one and a half day events where we invite researchers world-wide to give research seminars and meet with students. As a faculty member, I will continue these efforts in building interdisciplinary research communities.

Example Courses

I am qualified and excited to teach courses in the following areas:

Human-Computer Interaction Potential courses include introduction to HCI, interaction design studio, and user centered research methods. These are practice-oriented courses that teach the fundamentals of HCI and design. Topics include user experience research methods, prototyping, and usability testing.

Research Methods Courses Potential courses include PhD-level research methods courses, such as quantitative research methods and qualitative research methods. Quantitative research methods will cover topics such as experiment design, hypothesis testing, causal inference, etc. Qualitative research methods will cover topics such as observation, interview, grounded theory, etc.

Applied Machine Learning I plan to develop this course based on the CMU course taught by Prof. Carolyn Rose. The course will teach the concepts behind a variety of algorithms, such as decision trees, naive bayes classifiers, and emphasize on the rationale behind when to use them. Examples where such algorithms are successfully applied to address real life challenges will be given and discussed. Another emphasis will be on the process of effectively applying machine learning to problems of interest, from data scraping and cleaning to performing error analysis.

Web Development and Applications Topics would include modern web development technologies and frameworks. The course will be project-based, and students will be encouraged to design and develop web applications targeting real-world problems, with potential deployment at the end of the course.

Learning Technologies This graduate-level course will cover key topics in this area of research. It will examine bodies of research on intelligent tutoring systems, learning analytic methods, AR/VR applications, computer-supported collaborative learning, learnersourcing techniques, educational data mining techniques, educational games, etc. I plan to design two versions of this course. A basic version will focus on introducing and discussing the research trajectories in each of the areas. A studio version will be a project-based research seminar. Students will pick one area to do in-depth research with the goal of producing a ”work-in-progress” quality paper.