Maria Teleki

Teaching Statement

I have high expectations for my students. I believe they each have something individual and unique to contribute, and that my job is to help them make that contribution.

In this statement, I detail my teaching experience – earlier teaching, guest lecturing, and course development – and the courses I am prepared to teach, and finally my research mentoring experience.

Teaching

In my teaching, I am dedicated to not only teaching students the mechanics of engineering, but also the critical thinking skills around how and when to apply these skills.

Course Development. This past Spring, I TA'd for *CSCE 676: Data Mining (Graduate Level)*, where I helped redesign the course for the AI era. I developed assignments that required students to use coding agents to aggressively test their own implementations: in several homeworks, students had an agent write their version of a function and then validated it against an established library, learning to test their code from multiple angles and catch errors. I also wrote assignments in which students recorded themselves answering technical interview questions on video, training them to answer deliberately vague interview questions by stating their assumptions. Recently, I also helped shape *CSCE 289: GenAI for Engineers (Undergraduate Level)*, working through what students should learn about LLMs and lightweight coding to be effective in AI-driven engineering. I contributed to the progression from foundational concepts (tokens, probabilities, hallucinations) through hands-on API and pipeline work to critical evaluation of model outputs, so students leave able to use these tools competently and skeptically.

Guest Lectures. In Spring 2025, I guest lectured for CSCE 670 *Information Storage & Retrieval* (class size: 125) on Evaluation and Learning to Rank for search engines. I design my lectures knowing that this material informs the next generation of industry and academic leaders, and therefore set up the lectures to facilitate students' understanding (quotes shown are selected anonymous feedback; **full set of student evaluations available on request**):

> *In her slides, she clearly outlines a problem statement/goal early, with a clear explanation of why it matters, that being the relevancy of search results and its impacts- and also addresses why some algorithms work better than others, thus immediately following with the methodology she employs to address this.*

In my lectures, I engage students via quick Q&A, confidence polling (thumbs-up/-down), and small-group discussion. These techniques allow me to stay in-tune with the class's level of understanding, help students practice speaking up, and facilitate independent thinking among students. Students do well with these techniques:

> *She monitors our reactions frequently, e.g., thumbs-up... and it makes me feel more involved.*

I am able to encourage students to participate by regularly reminding them that *we are here to learn*, so it is acceptable to give me a wrong answer in class:

> *[She] took the load off of anyone to give "wrong" answers... even if someone was not right about an answer, she handled it very well.*

During these Q&A and discussions, I ask students questions like: *When is recall more important? When is precision more important?* And allow the class to discuss and share ideas. This type of discussion mimics the discussion that students will have as industry and research leaders.

Teaching Background. My teaching began before graduate school, designing and instructing Lego robotics and science classes for the YMCA Afterschool program in Williamson County, and before that as a robotics instructor and camp counselor at Camp Twin Lakes. The YMCA gave me some of the most valuable teacher training I've received: set expectations explicitly and remind students of them consistently, and most behavior problems never happen – a lesson that has held up just as well in the college classroom. As an undergraduate at Texas A&M, I worked as a peer teacher for introductory CS courses, holding weekly review sessions with notes I developed and assisting the TA at lab sections. Teaching introductory material showed me how many different angles students come at it from, and hammered home the importance of constantly checking for student understanding.

Courses

Core Courses. I am glad to teach core undergraduate courses, and would welcome large-enrollment sections in *Introduction to Programming*, *Data Structures and Algorithms*, *Probability and Statistics*, and *Machine Learning*.

Courses in My Area. At the undergraduate and graduate levels, I am prepared to teach *AI*, *NLP*, *Speech AI*, and *Computational Social Science*.

New Courses. As part of my CIRTL certification, I developed a sample syllabus for *Socially Robust AI* (available upon request). Students identify a social robustness failure, build the measurement for it first, and then build and compare techniques against that fixed measurement — delivering a new evaluation metric plus a method. The course is built for the interdisciplinary future of computing: cross-listed between computer science and information science, and drawing students from both. I would also be excited to develop courses in *AI Evaluation & Auditing* and *AI Governance*.

Mentoring

In my mentoring, I encourage students to think independently, and my role is to support them on their path to becoming strong collaborators and leaders. To this end, I have created highly flexible and personalizable structures which I iteratively update to support students in gaining and then applying their knowledge via project-based learning. These structures free my resources to provide timely and individualized feedback so I can scale the number of students I mentor, so that all my students have the opportunity to achieve those high expectations.
My philosophy centers around helping students develop and articulate their ideas; my job is to provide the scaffolding to allow them to grow. I provide structure to students in the following ways:

(i) *I hold weekly one-on-one meetings:* During our meeting, we brainstorm together: I ask students guiding questions to help them think through their roadblocks. Then we decide on what some reasonable next steps for the coming week are and document them in our shared Google Doc. This structure helps keep students moving through the project and not get stuck for too long in either the planning or implementation phase. It also allows students to take ownership of the project direction.

(ii) *I curate personalized projects:* I believe that we do our best work when we're excited about what we do. I regularly ask students, *What are you most excited about in our work right now?* Then, I help them figure out how they can study problems and methods that they are passionate about, so they can fully engage and help drive the project direction.

(iii) *I discuss positive, negative, and inconclusive results ahead of time:* Many students are quick to panic when an experiment shows results they did not expect. I preempt this panic by discussing with

students before running the experiment: *What does it mean if we get the results we expected? What if we get the opposite results? What if nothing happens/the results are inconclusive?* I find that having this discussion with students ahead of time calms students down and allows them to remain unbiased as they work through their project experiments. My aim is to keep students focused on learning.

(iv) *I created a repository of documents, resources, and materials:* These documents address parts of the research process which are common to all of my mentees – *How to Make a Poster/Presentation*, *Q&A Sessions*, *What to Expect at a Conference*, *How to Read/Write a Paper*, *Engineering Notes (Lab Server Use, Common Python Packages, etc.)*, a list of papers, datasets, and metrics that we find interesting, and more. Students are encouraged to add their knowledge to these documents. This repository frees my time because rather than re-teaching the same concepts to students, I can spend my time focusing on providing the individualized feedback that they need to move towards independence into the collaborator role.

(v) *I cultivate a culture of calm:* In all my interactions, I proactively create a safe space for students to ask questions, have a *rough* rough draft, question anything and everything, and discuss both struggles and successes. I regularly ask for feedback, so I can continually improve my mentoring, and efficiently move students forward in their research.

These structures have allowed me to grow a research program that empowers students to take ownership of their work and make research contributions. I now detail how I've developed my mentorship program over the course of my PhD.

The First Students (2023-2024). Soohwan Kim (UG) joined for course credit and co-authored [9]. Rohan Chaudhury (MS) followed as my first MS mentee, leading [3], and completing his thesis. Working this closely with two students showed me how much of what I was teaching was *exactly* the same each time – how to read a paper, how to structure a poster, how to get onto the lab server – and those repeated explanations became the first resources in my shared repository. My PhD advisor, James Caverlee, was very involved with these first two students, course correcting me and teaching me how to scope these projects.

Scaling the Program (2024-Now). With the shared repository in place, I started being able to support multiple students at once, and James took a step back to allow me to grow in this way. Sai Tejas Janjur (MS), Cong Wang (MS), Oliver Grabner (UG), Thomas Docog (UG), Ketan Verma (UG), Stephanie Birkelbach (UG), and Jason Kim (UG) worked in parallel on [10] and [11], with several also contributing to [8]. Each student owned a chunk of the work, and the weekly one-on-ones kept progress going. Sai won a 1st Place Award at Texas A&M's Student Research Week (SRW), and multiple students presented at both the Texas NLP Symposium and SRW; each student's authorship is detailed in my CV.

Taking Students from Contributor to Lead (2025-Now). I design each co-author project as an on-ramp to independent work, and students have taken it. Stephanie Birkelbach moved from contributing to [10] and [11] into leading SocialPulse [2], which she presented as an oral presentation at IC2S2 this year; this summer she completed an REU at the Rochester Institute of Technology with Ashique KhudaBukhsh, and is applying to PhD programs this cycle. Jason Kim followed the same path into leading PromptHelper [5], Anna Irmetova (UG) led PodChecker [4] for course credit, and Oluwadayo Bamgbelu (UG) led our work on demographic bias in LLM-based resume evaluation [1]. Junhyeon Kwon (MS) led *When Behavioral Traces Mislead* [6] while contributing to our real-time lecture comprehension tool [7]. Across twelve mentees, nearly all have co-authored, and six have led a paper as first author (four published, two in progress).

References

[1] Oluwadayo Bamgbelu, Maria Teleki*, Xiangjue Dong*, and James Caverlee. Detecting and Mitigating Demographic Bias in LLM-Based Resume Evaluation. In *In Preparation for ARR (Short)*, 2026.

[2] Stephanie Birkelbach, Maria Teleki, Nehul Bhatnagar, Peter Carragher, Xiangjue Dong, and James Caverlee. SocialPulse: An Open-Source Subreddit Sensemaking Toolkit. In *SocialLLM @ ICWSM '26*, 2026.

[3] Rohan Chaudhury, Maria Teleki, Xiangjue Dong, and James Caverlee. DACL: Disfluency Augmented Curriculum Learning for Fluent Text Generation. In *LREC-COLING '24*, 2024.

[4] Anna Irmetova, Haoran Liu, Maria Teleki, Peter Carragher, and James Caverlee. PodChecker: An Interpretable Fact-Checking Companion for Podcasts. In *MisD @ ICWSM '26*, 2026.

[5] Jason Kim, Maria Teleki, and James Caverlee. PromptHelper: A Prompt Recommender System for Encouraging Creativity in AI Chatbot Interactions. In *CHI '26 (Poster)*, 2026.

[6] Junhyeon Kwon, Jobin Varughese, Maria Teleki, and Mohsen Dorodchi. When Behavioral Traces Mislead: Evaluating Evidence Adequacy in Learning Analytics. In *In Prep. for LAK (Short)*, 2027.

[7] Amanda Lacy, Maria Teleki, Esau Hutcherson, Jobin Varughese, Jun Kwon, Frank Shipman, and Tracy Hammond. AI in the Loop: A Multimodal Assistant for Real-Time Lecture Comprehension. In *Learning at Scale Demo '26*, 2026.

[8] Maria Teleki, Vedangi Bengali, Xiangjue Dong, Sai Tejas Janjur, Haoran Liu, Tian Liu, Cong Wang, Ting Liu, Yin Zhang, Frank Shipman, and James Caverlee. A Survey on LLMs for Story Generation. In *EMNLP Findings '25*, 2025. Extended Version Is Under Submission at TIST.

[9] Maria Teleki, Xiangjue Dong, Soohwan Kim, and James Caverlee. Comparing ASR Systems in the Context of Speech Disfluencies. In *INTERSPEECH '24*, 2024.

[10] Maria Teleki, Sai Janjur, Haoran Liu, Oliver Grabner, Ketan Verma, Thomas Docog, Xiangjue Dong, Lingfeng Shi, Cong Wang, Stephanie Birkelbach, Jason Kim, Yin Zhang, and James Caverlee. Z-Scores: A Metric for Linguistically Assessing Disfluency Removal. In *ICASSP '26*, 2026.

[11] Maria Teleki, Sai Janjur, Haoran Liu, Oliver Grabner, Ketan Verma, Thomas Docog, Xiangjue Dong, Lingfeng Shi, Cong Wang, Stephanie Birkelbach, Jason Kim, Yin Zhang, Éva Székely, and James Caverlee. Conversational Speech Reveals Systematic Robustness Failures in LLMs. In *In Prep. for ARR*, 2026.