# Teaching Statement

**Mentoring:** As a research advisor, I aim to provide a welcoming, inspiring, collaborative space for young researchers to form and refine their independent ideas. My PhD research at Carnegie Mellon's Human-Computer Interaction Institute has exposed me to a diverse menagerie of interdisciplinary research approaches applied by the EdTech and Human-Computer Interaction (HCI) communities. I would be delighted to advise students who embody the spirit of interdisciplinarity extolled by research communities that I am proud to be a part of such as CHI (Human Factors in Computing Systems), AIED (AI in Education), Educational Data Mining (EDM), Cognitive Science (CogSci), Advances in Cognitive Systems (ACS) and many more. My efforts in publishing and reviewing for these communities have prepared me to advise students in carrying out plans for rapid publication within these areas.

As an advisor, I view student well-being and career support as my top priorities—a philosophy that I have applied to advising a total of 15 master's and undergraduate students, and co-advising four PhD students to date. My mentoring and research have complemented each other throughout my academic career. As I have matured as a researcher, I have relished learning to develop students' passions and ideas into novel research directions. To date, three of my advisees have successfully co-authored publications with me, and I expect to co-author several more this year as my mentoring load has grown in my role as a postdoctoral fellow.

**Classroom Instruction:** As a classroom instructor, my specialization in educational technology and the learning sciences makes me well-suited to deliver world-class instruction. As a fellow of Carnegie Mellon's Program for Interdisciplinary Education Research (PIER), I took **numerous classes in learning science** and spent over 80 hours observing teaching in practice in high school classrooms. Through PIER, I completed an interdisciplinary project with then Psychology PhD student Patience Stevens, in which we developed a mobile tutoring system that supported **early phonological awareness practice for kindergarteners** through automated, step-by-step adaptive feedback on invented spelling [1]. As an instructor, I am a firm believer in **backwards design**, the notion of first concretely defining learning objectives and then designing assessment and instruction to ensure student mastery. I am well-versed in best practices on providing effective feedback to students, leveraging educational technology effectively, providing metacognitive support, facilitating interactive social learning experiences, fostering inclusion and diversity in the classroom, and leveraging student feedback to refine course material.

During my PhD, I served as a teaching assistant for classes in **user interface programming** and design, and **e-learning design and theory**. A key takeaway from these experiences was facilitating students' transition between early-semester mastery-based learning to late-semester project-based learning, where they put course learning objectives into practice. From a course-design perspective, I see open-ended project work as an essential component of learning-by-doing for cultivating real-world expertise. In consideration of my future teaching, I look back fondly on the masterful executions of this project-based approach in UC Berkeley undergraduate courses such as CS 162: Operating Systems and System Programming, where our group projects involved implementing large parts of a Unix operating system. In leading and redesigning much of the curriculum for CMU's Software Structures for User Interfaces (SSUI) in the Spring of 2020, I applied this same mentality of setting a high bar for projects, and was impressed by what my students' were able to achieve in a short time by applying the principle from my curriculum to their own self-directed projects.

As an instructor, I am prepared to teach classes on a wide range of topics, including **educational technology**, **human-computer interaction**, **data science**, and **artificial intelligence**. I would be excited to develop courses on **AI in Education, Cognitive Systems, Computational Models of Learning,** and **Educational Data Mining.**

**Outreach and Community Building:** The route to scientific revolution is trailblazed by researchers who venture off the beaten path, which is why one of the most valuable things that an academic can do is to herd researchers venturing along similar new directions to build a community around a common set of novel ideas and ambitions. My outreach efforts have been heavily focused on building a community around cognitive systems research that focuses on the study and simulation of human learning.

The idea of studying human learning through the development of AI is an old one, but *not yet* a revolutionary one. The remarkable successes in this space have yet to build a following similar to applications of data-driven machine learning and generative AI, but are nonetheless natural solutions to the many follies of popular AI approaches especially when applied to education—their inaccuracy, lack of interpretability, data-hungriness and the pernicious way in which they have redefined the notion of a *model* in their own image; as just tools for prediction instead of scientific analysis. Many have proposed neurosymbolic AI as a solution to these faults. Yet, most proposals along these lines have a highly simplistic view of the symbolic side of the equation, like, for instance, simply augmenting generative AI by drawing information from documents or databases or calling out to ready-made external programs. Average AI and EdTech researchers resign themselves to a limited set of tools, not because they are the right tools, but because those tools have broad successes and sustained momentum. Few are aware of the many cases where symbolic approaches to machine learning and AI have eclipsed artificial neural networks as both practical and scientific tools. For instance, had I resigned myself solely to gradient-descent-based learning mechanisms like deep reinforcement learning, as many AI researchers do, my simulated learners could never have matched human learning rates, nor would it be possible to investigate their changes in knowledge directly and develop theories of human learning around a well-defined theory of multiple learning mechanisms.

Over the last 6 years, as a leader of the Computational Models of Learning track at CMU's LearnLab summer school, I have instructed almost 100 learning engineers in how to use my simulated learners to investigate different instructional design decisions in intelligent tutoring systems. Paradoxically, interest in this workshop has grown with the rise of large language models, since my simulated learners provide a path for learning engineers to do what they could not do with generative AI alone: simulate student learning instead of just student behavior. This year at Advances in Cognitive Systems, I encountered many similarly ambitious projects of research that successfully simulated various elements of human cognition with a combination of symbolic and generative AI. This gathering reminded me a bit of the state of neural network research in the early 2000s: many small successes distributed across several bespoke projects of research with unique special-purpose tools. A great deal could be collectively accomplished within this community if we had a common standardized toolset; something analogous to PyTorch for cognitive systems. Toward developing such a toolset, I am currently packaging the efficient algorithms that I developed for building my most recent simulated learner AI2T into a C++ compiled Python extension called the Cognitive Rule Engine (CRE). Soon, I intend to release AI2T and CRE as tools for a broader audience; AI2T for learning engineers, and CRE to bootstrap other researchers' unique ideas about how to simulate human learning and cognition. In terms of community building, I intend to organize workshops at AIED and other EdTech conferences on simulated learners and host symposia at venues like AAAI.

## References


1. **Weitekamp, D.**, & Stevens, P. (2022). A Mobile Invented Spelling Tutoring System. International Conference on Artificial Intelligence in Education, 492–496. Springer International Publishing Cham.