I. Research and Creativity

Artificial Intelligence (AI) technologies are on course to transform society, with the potential to revolutionize how people learn and work. In education, AI offers the promise of personalized instruction at scale—giving everyone access to *tutors* that adapt to their needs to enhance learning. However, cost-effectively creating AI tutors that are reliable and pedagogically effective is difficult.

This challenge is especially relevant with the recent rise of large language models (LLMs) such as ChatGPT. While these models often produce fluent responses across tasks, they have hidden gaps in their knowledge and as a result are prone to generating incorrect or misleading information. For example, a recent study from my lab evaluating ChatGPT models on math tutoring found that although they generated the correct final answer for 85.5% of tested college algebra problems, only 56.6% of their tutoring dialogues were entirely correct [B2.37]. In contrast, intelligent tutoring systems are reliable and effective,[1] but are costly to develop and customize to end user needs.

My research addresses this issue by adopting a ***cognitive systems*** approach to ***AI in education***. Cognitive systems research aims to "*explain the mind in computational terms and to reproduce the entire range of human cognitive abilities in computational artifacts*".[2] I leverage this approach to develop AI technologies that effectively support humans. These technologies are then deployed in real-world settings to generate meaningful, measurable impact. In a virtuous feedback loop, data collected from these deployments informs and advances my cognitive systems research—leading to improved AI capabilities, refined computational models, and ultimately, more effective technologies for future deployment (see Figure 1). My lab applies this methodology to better understand how people teach and learn, and to build machines that can teach and learn like people do. Our work draws on and contributes to several fields including cognitive science, machine learning, human-computer interaction, and the learning sciences. By integrating across these areas, my lab has created AI tutors and deployed them in over **300 courses** serving more than **5,000 learners**, along with novel teachable agents that let end users create and personalize AI systems.

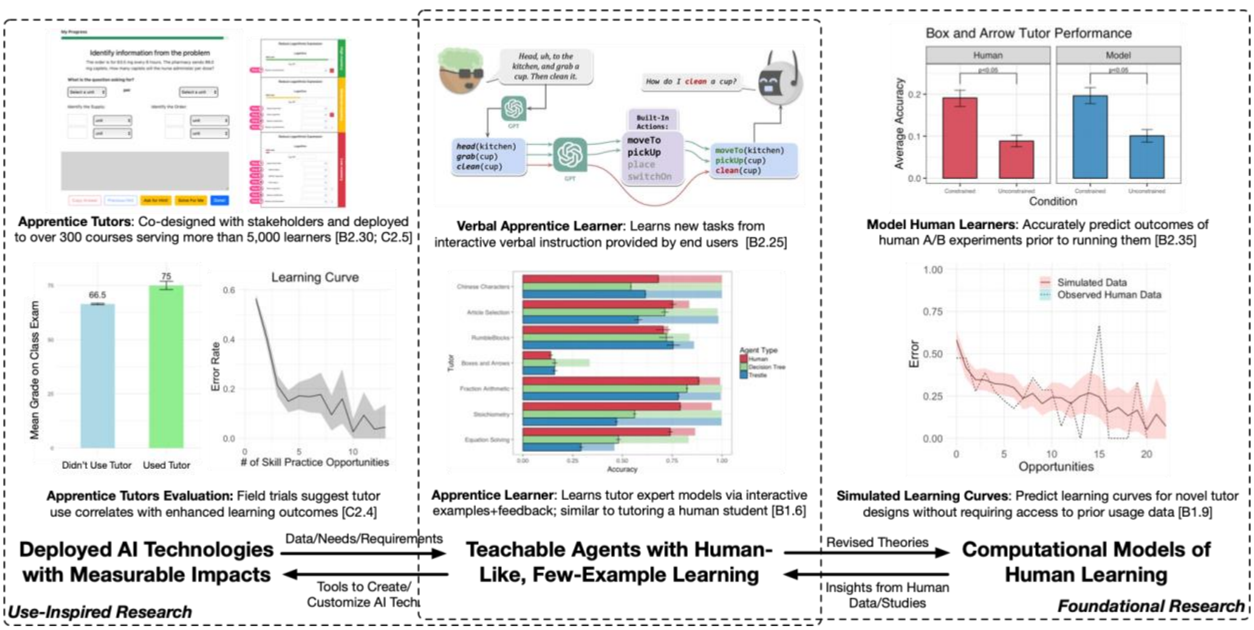

Figure 1. The three elements of my research program, how they relate, and select examples of work in each.

[1] Pane, J. F., Griffin, B. A., McCaffrey, D. F., & Karam, R. (2014). Effectiveness of cognitive tutor algebra I at scale. *Educational Evaluation and Policy Analysis*, *36*(2), 127-144.
[2] Langley, P. (2012). The cognitive systems paradigm. *Advances in Cognitive Systems*, *1*(1), 3-13.

Deployed AI Technologies with Measurable Impacts

My research places an emphasis on translating AI technologies from the lab into the world, so they can have broader impacts and ground my lab's human-centered AI work. For example, through the NSF AI ALOE Institute, my students and I have engaged in co-design with stakeholders at the Technical College System of Georgia (TCSG) to develop and refine more than 46 AI tutors for mathematics and nursing education that align with instructor and student needs (see upper left image in Figure 1). This process has produced new intelligent tutoring system techniques, such as hierarchical task network-based tutors that capture the compositional nature of skills and can dynamically adapt how much scaffolding they provide based on a learner's mastery [B2.30]. Deploying and maintaining these tutors for real-world use takes substantial time and resources beyond typical academic work, but our efforts have produced meaningful impacts. **Our deployment data suggest that students who use our tutors achieve higher course grades than those who do not and that improved error rates and test scores are correlated with greater tutor use** (see bottom left in Figure 1) [C2.4]. Now that we have tutoring systems successfully deployed, we are currently running A/B experiments through our tutor platform to causally evaluate the effects of AI tutor capabilities on learning outcomes.

The data collected from our AI tutors has also enabled other lines of research that inform how researchers, educators, and student use AI in the classroom. For example, many people are turning to LLMs for tutoring support, but it is still unclear how accurate or reliable they are for this purpose. To systematically evaluate LLMs for tutoring, my students and I created new benchmarking methods that leverage tutors that have been deployed in actual classrooms and the data produced by learners using these tutors. Our methodology evaluates how reliable and accurate LLMs are at tutoring learners on the exact tasks intelligent tutors currently support in the classroom. Our findings suggest that while LLMs do reasonably well at generating final answers to tutor problems (around 85% accuracy), they have serious shortcomings—we found that around half of the LLM tutoring dialogs contain errors on intermediate steps [B2.37] and that LLMs struggle to provide accurate correctness feedback across the full range of student inputs encountered in deployed tutors [B2.38]. This work provides new ways to empirically measures the performance of LLMs at tutoring, so stakeholders can better evaluate the risk-benefit tradeoffs of using this technology in the classroom.

Teachable Agents with Human-Like, Few-Example Learning

To increase access to reliable AI tutoring support, my lab has pioneered new interactive machine learning approaches that let teachers create and personalize AI tutors for their classes via natural teaching interactions—similar to how they would teach one of their students (see Figure 2).

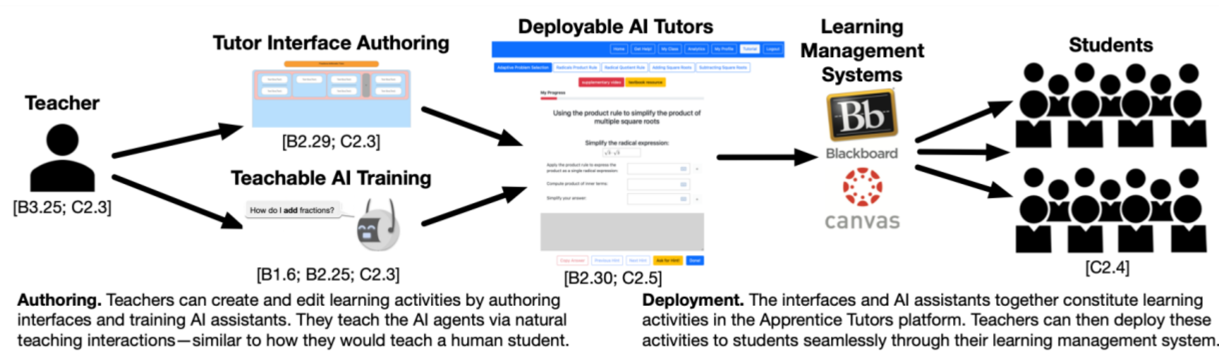

Figure 2. How teachable agents situate within my lab's AI tutor deployment pipeline. Recent publications on different elements of the overall pipeline are also highlighted.

Within our pipeline, teachers start by creating tutor interfaces to scaffold the learner's problem-solving process. Most teachers are not interface designers, so my lab has developed innovative new tools to make authoring tutor interfaces easier for teachers. For example, we developed a new direct manipulation approach that lets teachers create an interface in less than 10 minutes [C2.3]. In subsequent work, we invented a new technique that uses generative AI to make interface authoring faster and easier for teachers [B2.29]. My lab's tools are designed to balance autonomy with control—intermixing generative AI authoring with direct manipulation to get the best of both. We have discovered that while generative AI is unreliable for directly tutoring students, it has under-explored potential for helping educators create and personalize instructional materials.

After creating an interface, teachers train an AI agent to solve problems in it much like they would teach a student—via tutoring. Once trained, the agent can be deployed alongside the interface to provide reliable, pedagogically aligned support to students at scale. A central challenge for teachable agents is learning effectively from the limited data a single instructor can provide. To address this, my lab has discovered new ways for people to teach AI agents with minimal guidance. For example, my lab created AI agents that learn and update symbolic task knowledge from interactive examples and feedback [B2.7; B1.6; C2.3] as well as from verbal instructions [B2.25]. These approaches are substantially more data efficient than alternatives like reinforcement learning [B2.19], often enabling agents to learn new behaviors from just a single example. While my earlier work evaluated teachable agents with technical users [B2.7; B2.14; B1.6], my recent work emphasizes testing with end users, showing they can successfully teach agents after only a few minutes of working with them [B2.25; C2.3]. My work uniquely combines generative AI approaches, which we use to make interaction easy and natural, with symbolic AI approaches, which provide efficient, reliable, and interpretable learning and reasoning capabilities. With this combination, **my lab's teachable agents are easy to use, quick to train, and can reliably deliver accurate tutoring without the hallucinations that typically plague LLMs.**

Computational Models of Human Learning

To power my lab's teachable agents and guide the development of AI tutors, I build computational models of human learning that explain how people engage in efficient, continual, and few-example learning—and that replicate these capabilities in AI systems. I am an innovator in this space, having developed several models with human-like learning capabilities and demonstrated multiple novel, practical applications (see right pane in Figure 1). For example, **I introduced a new approach that uses computational models to predict student learning curves without requiring student data—a theory-driven, data-free method** [B2.13]. This work was recognized with an Exemplary Paper Award at the Educational Data Mining conference. I have also shown that computational models can predict the outcomes of hypothetical and counterfactual human experiments [B1.9; B2.35], opening new possibilities for instructional design. This high-impact work has the potential to transform the tutor design process by providing instructional designers with "simulated students"—enabling them to efficiently test and refine tutor designs before real-world deployment. Research on simulated students is an emerging area of interest in the AI in Education community, and I continue to be a trendsetter in this space. Most recently, my lab released a new testbed for evaluating simulated students using tutoring systems and human learner data [B2.38].

Broader Applications of My Research

Although my lab primarily focuses on AI in education, our research has attracted significant interest and recognition in broader domains. For example, DARPA funded my lab to apply our cognitive systems approach to medical diagnosis. In this project, we developed an innovative, patent-pending method that leverages both doctor insights and human-like sparse coding to achieve state-of-the-art diagnostic performance with very limited training data [B2.23; B2.31]. In an independent evaluation by the FDA, this technique outperformed all other DARPA performers on the project and helped shape the FDA's emerging perspective on evaluating AI models for medical applications. Building on this work, my lab is developing new models of human-like concept formation that support data-efficient, continual learning without catastrophic forgetting [B2.17; B2.26; B2.27; B2.28]. In an ongoing project funded by the Army Research Lab (ARL), we are leveraging teachable agents to facilitate more effective human-AI teaming on cooperative tasks [B2.22; B3.23; B3.26; B2.32]. These efforts further demonstrate the versatility and real-world impact of my lab's research.

II. Teaching and Educational Contributions

I view teaching and mentoring as complementary to research and have prioritized both. Beyond receiving training in learning science, I have spent over 80 hours observing classrooms and have designed/taught courses in AI, HCI, and data science. Core pillars of my teaching philosophy include creating learning environments that center intellectual diversity, critical thinking, and respect, leveraging feedback to continuously improve, and personalizing my mentoring to each student.

Creating learning environments that center intellectual diversity, critical thinking, and respect. In my classes, I intentionally foster a diverse range of ideas and perspectives, encouraging my students to critically reflect on, discuss, and challenge these views. When teaching my Human-AI Interaction course (taught at GT in Spring '24 and currently ongoing in Spring '25), I intentionally included a broad range of perspectives on each topic (technological, social, ethical, and cultural) and used class time for collaborative activities that promote deep engagement. I believe my efforts were successful, with several students sharing encouraging comments on my course surveys:

- *"I think the course was exceptional because the quality of the content that was discussed and reflected upon was very well thought-out with respect to establishing continuity with prior discussions from preceding readings and thought-provoking in the sense that most of them just begged students to take a stance on radical ideas."*
- *"[Chris] pushes his students to consider difficult topics that don't always have a straightforward answer, and he encourages us to speak up on these topics"*
- *"The ethical discussions and listening to all the different perspectives that might not have thought about otherwise creates a bigger empathy and better understanding in AI, especially at a University that is known for its AI program", and*
- *"I very much enjoyed HAI and I was more engaged in this class than almost any other."*

These qualitative indicators supplement my quantitative indicators (e.g., instructor effectiveness of 4.95/5 in Spring '24) to demonstrate the value of my unique approach. I have also received external recognition as an educator, including the 2022 AAAI/ACM SIGAI New and Future AI Educator award for my ideas on interdisciplinary approaches to AI education.

Leveraging feedback to continuously improve my instruction. I believe that no course design is perfect and continuous revision is necessary to stay relevant—particularly in AI. I adopt a learning engineering philosophy that uses data to guide re-design. After teaching knowledge-based AI in Fall '22, I received feedback that students needed a clearer mapping from theoretical concepts to practical applications. I redesigned core elements of the course, including restructuring topics, redesigning the assignments, and creating a new course project. In response, my CIOS ratings for "course effectiveness" improved from 4.18 to 4.57 between Fall '22 and Fall '23. In the most recent round of feedback, I learned that students wanted me to better show the relevance of course concepts to the current (data-driven) AI discourse. This past summer, I redesigned the project to center on the Abstraction and Reasoning Corpus (ARC)—an AI benchmark that focuses on Artificial General Intelligence (see https://arcprize.org). This benchmark is designed to require reasoning in low-data scenarios, where knowledge-based AI is particularly relevant. The course project aligns with a public Kaggle competition based on ARC that offers $1 million in prize funding for competitive approaches and thoughtful papers, and the project is designed to facilitate students in competitively participating (if they choose). By linking course ideas to a widely publicized competition, I aimed to better demonstrate the importance and relevance of knowledge-based AI and to encourage students to become active participants in the public AI discourse. In response to these changes, my CIOS ratings for "course effectiveness" further improved to 4.72 in Fall '24.

Personalizing my mentoring to each student's unique needs. I mentor students from diverse backgrounds, including AI, ML, HCI, and the cognitive and learning sciences. Every mentoring relationship is different, so I always start by working to better understand each student's goals and experiences, which makes it possible for me to personalize my guidance to their needs. Early on, I work closely with students and gradually shift to a more supporting role as their research competencies grow. Using this apprenticeship approach, my students are able to publish quickly, learn while doing, and grow as independent researchers. I have also created a culture of co-mentorship in my lab, with more senior students providing advice and mentorship to more junior ones. This has enabled me to expand my individual mentorship capacity to include 39 students since joining Georgia Tech in Fall 2022 (8 PhD, 1 visiting PhD, 2 postdoctoral, 19 MS, and 9 undergrad). There are several recent indicators of my mentorship success. For example, my first two PhD students graduated this past year, several of my students have received awards, and many of my MS and undergraduate students have gone on to start PhD programs.

III. Service

I actively engage in high-impact service both internally and externally. In 2022 and 2023, I co-organized Interactive Computing's Graduating PhD student mentoring program (with Josiah Hester and HyunJoo Oh). I worked with my colleagues to develop a mentorship curriculum, hosted monthly meetings over the academic years, and provided 1:1 consultation to help students prepare for their next career steps. This year, I serve on the school's faculty hiring committee. I also support multiple faculty area committees (AI, Cog/Learning Sci, HCC, and HCI), through which I have updated the AI specialization reading list for HCC, supported writing and evaluation of 18 qualifier exams, served on 9 dissertation committees, and assisted in admissions across these areas. I have also been a primary contributor to the NSF AI ALOE Institute, where I serve as a member of the executive committee and co-organizer (with Scott Crossley) of the Foundational AI working group.

I also provide leadership and service within my professional communities. I have developed and taught summer schools and tutorials on computational models of learning at CMU and AAAI, organized several human-AI teaming workshops in partnership with ARL, and organized the first AAAI Symposium on Human-Like Learning at Stanford, among several other activities. I am also serving as the general chair for a conference in my area of specialization (Advances in Cognitive Systems), which I will host at Georgia Tech in Fall 2025. Lastly, I annually serve as a senior reviewer for several conferences and journals as well as a grant reviewer at NSF and IES.

In conclusion, I aim to have a transformative impact on the world in the areas of cognitive systems and AI in education. In the short time I have been at Georgia Tech, I have demonstrated excellence as both an educator and researcher, bringing a unique perspective to AI that crosses disciplinary boundaries. Additionally, I have been a mentor and leader, fostering the next generation of scientists and actively contributing to both the institute and the broader community.