THE AI-ALOE SPOTLIGHT

SPRING 2023

In This Issue

FROM THE EXECUTIVE DIRECTOR

The National AI Institute for Adult Learning and Online Education (AI-ALOE) is pleased to introduce its quarterly research newsletter: The AI-ALOE Spotlight. Our Institute will publish a newsletter every few months. It will include updates on the latest research happening at AI-ALOE. Our research covers various topics such as Foundational AI and Use-Inspired AI, adult learning and online education, participatory design for responsible AI, and technology infrastructure for deploying using AI to support adult online learning.

An example of our work involves developing a technology infrastructure for deploying AI cognitive assistants to enhance online teaching and learning. We gather and analyze data on learning, which is then used to inform both the AI cognitive assistants and the teachers and learners. Furthermore, in our Foundational AI work, we are tackling critical research areas related to human-AI collaboration, including the mutual theory of mind between humans and AI agents, self-explanatory AI agents, teachable AI agents, information visualization, and personalization of learning for adult online learners.

Our Institute is also actively engaged in collaboration and outreach to build a nexus of connections and activities around AI for learning and education. Case in point, we organized a virtual symposium surrounding this topic in November 2022; we also participated in a forum on the topic organized by Digital Promise in Mountain View, California in December 2022, and in April 2023, AI-ALOE will lead a panel that brings the three AI Institutes for learning and education together at the annual meeting of the American Educational Research Association (AERA) in Chicago. I will also add that, in May 2023, AI-ALOE will help bring the ethics groups of all 19 AI institutes funded by the National Science Foundation (NSF) and the National Institute for Food and Agriculture (USDA-NIFA) for a major conference at Georgia Tech. We are also seeking collaboration with Minority Serving Institutions, or MSIs, as part of the NSF's Expand AI program, and this newsletter is another example of our commitment to outreach.

Each issue of the AI-ALOE Spotlight will highlight the research of one faculty member and one student researcher who are engaged with AI-ALOE. In this issue, we shine a spotlight on Christopher MacLellan, an assistant professor in the School of Interactive Computing at Georgia Institute of Technology, and Langdon Holmes, a Ph.D. student in the Department of Psychology and Human Development at Vanderbilt University. We value your feedback and hope you enjoy the newsletter.

- Ashok Goel

Get Involved

Connect, subscribe, and follow AI-ALOE!

AI-ALOE Annual Retreat





The AI-ALOE Institute held its second annual Retreat at Georgia Tech, February 8-10, 2023. This was a three-day, full-team, mostly in-person event attended by more than forty team members, many of whom traveled to the city of Atlanta specifically for the event. The Retreat started with the AI-ALOE Executive Council meeting, on the first day, where members discussed the goals of the AI-ALOE retreat and overall goals of the Institute.

The next morning, a dozen AI-ALOE researchers gave short talks summarizing the progress and outcomes of their research projects, followed by a summary of the evaluation of the program by Dr. Aileen Reid. In the early afternoon, Dr. Cayanna Good, the Assistant Commissioner for Adult Education at the Technical College System of Georgia, gave us a talk that reminded us about why we do what we do: to improve the human condition. On behalf of the AI-ALOE's Learning and Education group, Dr. Chris Dede then discussed who the adult learner is for AI-ALOE and laid out a research plan for addressing adult learning. Later in the afternoon, Dr. Ashok Goel described an integrative framework for all AI-ALOE research projects and presented a logic model for the Institute.

On the third day, Dr. Dede, representing the AI-ALOE program review committee, shared the findings and recommendations from a months-long rigorous review of all AI-ALOE research projects. Dr. Goel then presented human-AI collaboration as an integrative framework for AI-ALOE research on foundational AI. In the afternoon, Dr. Goel, and AI-ALOE research scientists Harsh Sikka and Vrinda Nandan, presented plans for further development of the AI-ALOE technology infrastructure in the second year. Finally, Dr. Chaohua Ou described AI-ALOE's collaboration and outreach activities, focusing on developing programs for teaching fellows, research fellows, as well as undergraduate research fellows.

The presentations at the Retreat were interspersed with both small-group discussions and full-team discussions organized by Dr. Myk Garn. The discussions were always very energetic, and sometimes quite intense. At the end of the Retreat, the AI-ALOE team had a new shared understanding of its vision and mission, as well as a renewed commitment to realizing the vision and accomplishing the mission.

AI-ALOE Research & Partners

AI-ALOE is funded by the National Science Foundation (NSF) and Accenture through NSF. Led by the Georgia Institute of Technology, AI-ALOE aims to develop AI technologies to transform and advance online adult learning in **effectiveness, efficiency, access, scale, and personalization**. Currently, AI-ALOE has five research strands, and the figure below illustrates its main research themes, AI technology tools, methods, testbeds, and products:

AI-ALOE Research

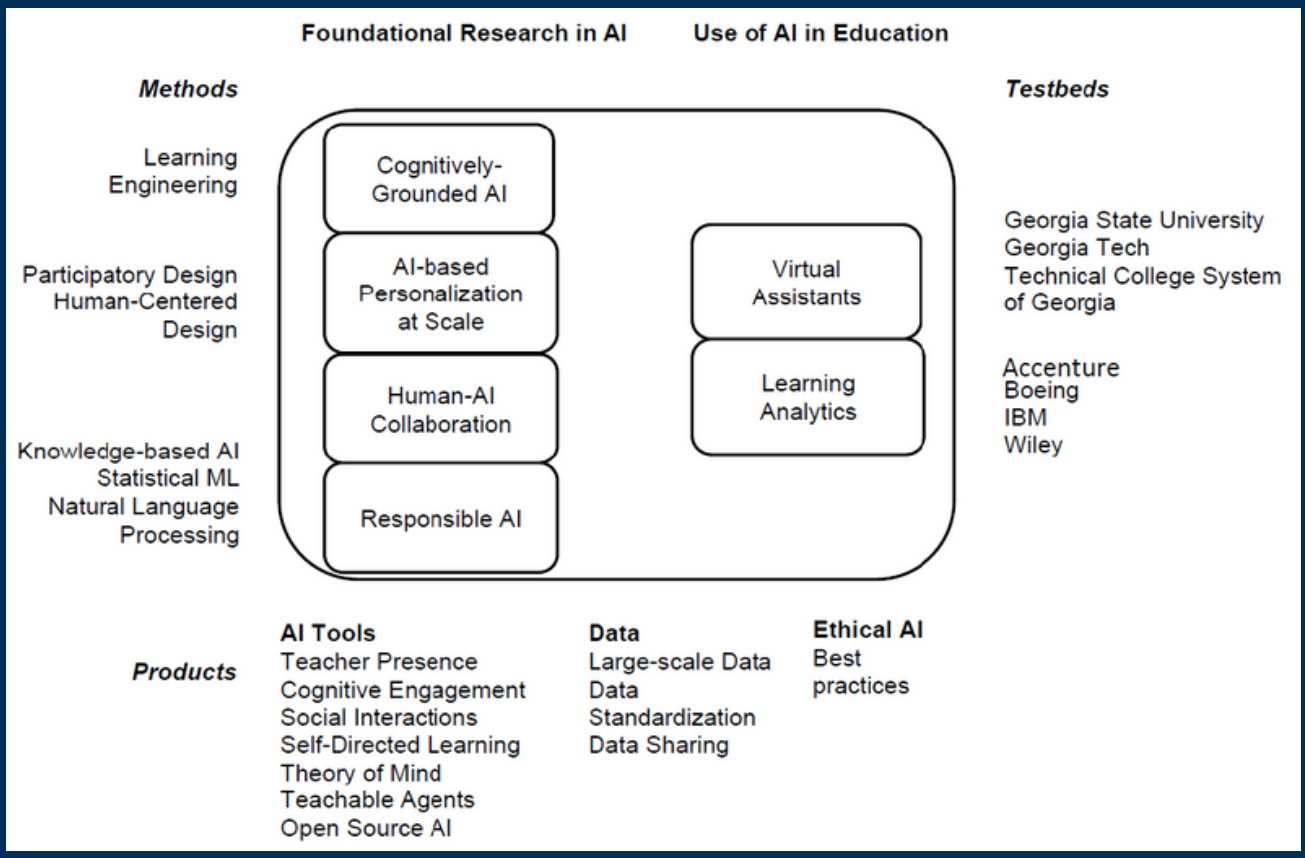

AI-ALOE Organizations & Industrial Partners

Georgia Tech · 1EDTECH Power Learner Potential · Georgia State University · Harvard University · TCSG Technical College System of Georgia · UNC Greensboro · Vanderbilt University · accenture · Boeing · IBM · Wiley

Technology Infrastructure

At AI-ALOE, our research is derived from two positive, data-driven feedback loops in human learning: (1) a teacher-centered loop, and (2) a learner-centered loop. As individual research groups, we build AI-based technologies that could fulfill sections of these loops. The challenge for AI-ALOE as a collective is how we can build an AI-based solution that realizes the full extent of these learning loops. To address this challenge, we are architecting an AI-based platform that allows multiple AI agents to jointly support our users' learning goals. Here is the **AI-ALOE technology infrastructure architecture**.

Data Pipeline Milestones

1. I Iterated on the production version of the data infrastructure using AWS, adding systematic and secure upload and download processes.
2. Collected Data from five AI technologies, deployed in dozens of classes at two different university systems, Georgia Tech and the Technical College System of Georgia (TCSG).
3. Collected student discussion data and other learning environment analytics.
4. Implemented a rule-based anonymization approach with 95% personally identifiable information (PII) removal success.
5. Validated and tested an obfuscation strategy, and designed and developed a Goal 1.0 version of the anonymization pipeline that includes Named Entity Recognition (NER) based recognition.
6. Developed documentation to support the onboarding of AI-ALOE researchers, university partners, and other members of the ecosystem in publishing and understanding data.
7. Updated and created artifacts to understand the status of various data in our pipeline through our data auditing efforts.
8. identified emerging considerations in analytics for personalization, including different scales of time and scope in the learning environment, different target users, personalization of heterogeneous AI technologies, ecosystems of collaborative technologies, and the scale of deployments and access.

User-Centered AI Platform Design

AI-ALOE is dedicated to creating personalized and seamless user experiences by developing user-centered AI platforms that help users better understand and engage with learning data. The development includes three key aspects:

1. mapping user journeys;
2. building effective dashboards; and
3. analyzing data.

By mapping out User Journeys, we gain a deeper understanding of the user's needs and desires. By understanding user needs, we can design effective dashboards that provide users with clear and concise data visualization and easy access to key information.

We have constructed a user journey for a fictional learner, Jonna Garcia. We also developed dashboard prototypes for SAMI and Apprentice Tutors that help researchers like Dr. Ayana Jones (a fictional user in AI-ALOE's researcher journey) better understand learner behavior.

Please click the images below to learn more about the learner journey and the dashboard prototypes.

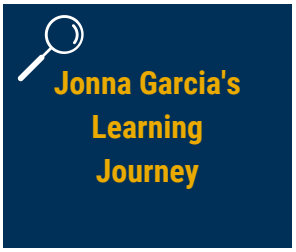

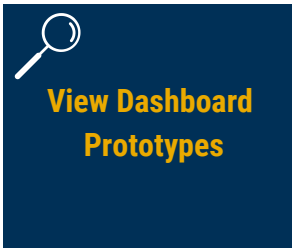

Apprentice Tutors

Research has proven that students benefit from 1:1 in-person tutoring, but it is challenging to provide this support to all students when they need it, especially in under-resourced institutions. To address this challenge, the **Teachable AI Lab** at Georgia Tech has made significant strides in advancing the development and deployment of intelligent tutors through their Apprentice Tutors platform and novel Teachable AI capabilities. Led by **Dr. Christopher J. MacLellan**, the lab's research has been focused on creating more personalized learning experiences for students at scale with Artificial Intelligence.

Apprentice Tutors

The Apprentice Tutors platform enables teachers to deploy intelligent tutors through learning management systems, such as Blackboard and Canvas. The platform utilizes AI techniques to analyze student data and provide personalized tutoring support to students as they are solving problems. Recently, the lab has been working with teachers to identify and develop new tutors for the Technical College System of Georgia's (TCSG) college algebra course. They have added an improved teacher's dashboard to the platform, which gives teachers insight into student learning on the platform.

The lab is also working on extending the platform's support for the Learning Tools Interoperability (LTI) standard, which enables seamless integration with Blackboard. The current system lets teachers select a tutor and deploy it through Blackboard with the click of a button. The lab is currently working on new functionality that will let the Apprentice Tutors platform report grades on tutored activities back to Blackboard, so teachers can assign practice with tutors as a graded activity.

As tutors have been deployed into TCSG classrooms, the Apprentice Tutors platform has been collecting data about how students learn. The lab is leveraging this data to create new AI approaches for tracking students' knowledge of skills. The platform currently utilizes an AI approach called Bayesian Knowledge Tracing, which estimates what students know and enables personalized next problem recommendations for students. The lab is exploring extensions to this approach as well as the translation of emerging approaches, such as Deep Knowledge Tracing, to enable enhanced personalization.

Apprentice Agents

The ultimate goal of the Teachable AI Lab is to create teachable AI agents that let teachers author and customize intelligent tutors by teaching a virtual agent like they would a human. Through teaching, the agent will learn how to perform the target task and can provide personalized one-on-one tutoring support to all the students in the class—like a digital teaching assistant. The lab is developing fundamental advances in interactive machine teaching to support this use case. They have been designing a new teachable agent that uses hierarchical task networks as their underlying representation, enabling the construction of new skills composed of previously learned skills. This new representation will enable better transfer of learned skills from one situation to another and will better support multi-modal teaching interactions that include a combination of demonstrations, feedback, and language-based goal guidance.

The focus of the lab in this quarter has been on creating an agent that can use hand-built hierarchical task network representations to solve problems and to support tutoring guidance. Moving forward, the lab is starting to explore adding different machine learning capabilities to the agent to acquire new knowledge from multi-modal instruction provided by a human trainer.

SMART

The **AI2 Research Lab** at Georgia State University is led by **Dr. Min Kyu Kim**. who works along with graduate assistants, Jinho Kim and Yoojin Bae, on the research and development of **SMART** (Student Mental Model Analyzer for Research and Teaching), an AI-powered system for formative assessment and feedback. Their two primary research goals are: (1) to explore data analytics methods and models that improve SMART formative assessment and feedback and analyze discussion posts, and (2) to investigate the impact of Year 1 SMART implementation in the English and Biology classes at the Technical College System of Georgia (TCSG).

To improve data analysis models for SMART and learning analytics for adult learners' online discussion, the research team conducted several projects, including:

- Validating and exploring SMART analytics in terms of measurable learning progress models. Specifically, the team compared SMART measures with human-rated scores and found acceptable correlations between AI and human-rated scores. The findings suggested that AI-driven assessment could provide reliable formative feedback. Additionally, they used SMART measures to create a learning progress model that determines students' conceptual change when conceptualizing complex problems.
- Leveraging a large language model, BERT, to develop machine learning models that automatically determine the levels of learners' cognitive presence in their discussion posts.
- Exploring data analytics techniques such as sequence mining and Hidden Markov Model (HMM) that can cluster learners' hidden states of cognitive engagement and their trajectories represented in longitudinal log data (i.e., a series of online discussions).
- Proposing a comprehensive AI literacy framework from a developmental perspective and applying the framework to examine how ALOE instructors developed their AI literacy.

Results and Impacts of Year 1 SMART Implementation

In December 2022, the research team completed Year 1 data collection from 11 sections of the English course at Columbus Technical College and 3 sections of the Biology course at Chattahoochee Technical College, where SMART benefited a total of 316 students. Leveraging multi-source data, including student surveys, SMART measures, instructor interviews, student interviews, and LMS data, the team investigated how the learners' experience on SMART influenced their concept learning, technology acceptance, discussion quality, metacognition, perceived value of Community of Inquiry, and final scores.

The results showed that learners in both courses significantly improved concept learning. The learner's experience with SMART positively influenced their perceived usefulness of AI-enabled learning in the Biology course. Learners in the English course demonstrated higher final scores when they conceptualized an appropriate model of the reading materials via SMART. The research team submitted two proposals on the research findings to the 2023 AECT International Conference.

During Year 2, the research team intends to collaborate with other ALOE researchers to enhance personalized learning and communication between humans and AI on SMART. They plan to work with Dr. Endert and his team at the Visual Analytics Lab at GaTech to create data visualization features on SMART. They will also work with Dr. Scott Crossley and his team at Vanderbilt to integrate a keyword identification tool to improve the expertise modeling on SMART.

Click to watch demo of SMART

Jill Watson

A common challenge in online education is the lack of teacher presence: online students lack the kind of access to teachers (and teaching assistants) that face-to-face learners often take for granted in physical classrooms. This is one of the main reasons for the low engagement and retention in online classes. Jill Watson is a virtual teaching assistant developed by the **Design & Intelligence Lab** led by **Professor Ashok Goel** at Georgia Tech. Jill Watson uses AI, machine learning, and natural language processing to enhance teacher presence in online classrooms: Jill answers student questions concerning course syllabus, policies, and schedules anytime, anyplace, even when no teacher or teaching assistant is available. We are now able to create and deploy Jill Watson in a course in four hours.

Jill Watson is currently deployed as a public Q&A agent on a class discussion forum, e.g., Edstem, and as a private Q&A agent on Canvas LTI (Learning Tools Interoperability). About 2,500 students are currently exposed to JW per semester in five computer science classes at Georgia Tech.

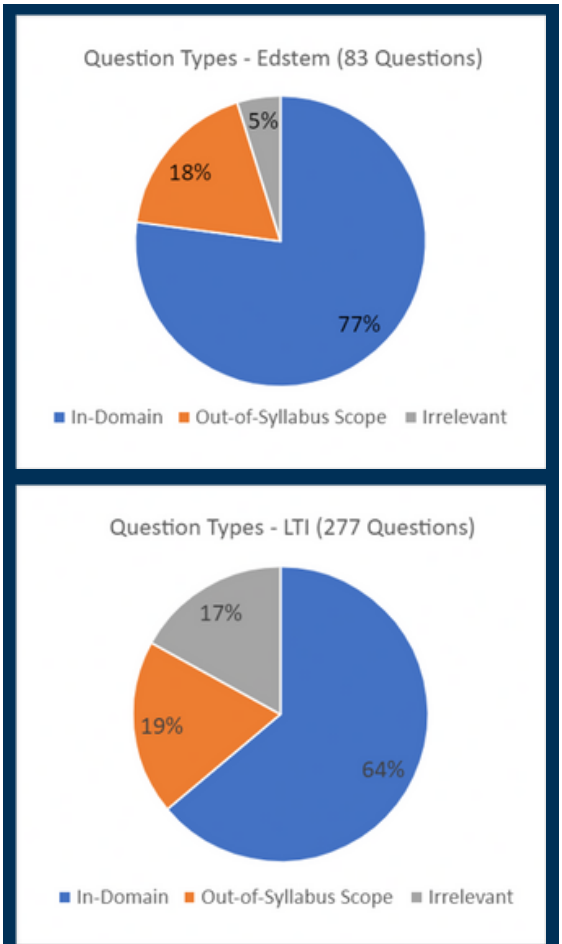

Types of questions asked on the public Edstem forum (top) and the private Canvas LTI (bottom)

In the public deployment of Jill Watson students can see questions posed by other students along with Jill Watson's responses. In the private deployment, students engage one-on-one with Jill Watson through a chatbot. One of the research questions that we have tried to answer is whether there is a difference in how students interact with the public agent in relation to a private agent. The figure on the left shows the preliminary results of this inquiry.

We put questions into three categories: (1) In-domain questions are related to the course syllabus and policy, e.g., "When is assignment 1 due?"; (2) Out-of-syllabus questions are relevant to the course but out of Jill Watson's domain, e.g., "When will assignment grades be returned?"; and (3) Irrelevant questions that lie outside of Jill Watson's domain, e.g. "Hi Jill, can you pass the Turing test?". Jill Watson is expected to answer only those in the first category.

The results show that students asked more questions in the private setting (277 questions) than in the public setting (83). This may indicate that students feel less intimidated asking questions in a one-on-one setting or/and they find asking questions in a dedicated chatbot more convenient. In addition, in the absence of the peer gaze in the private setting, students ask a lot more irrelevant questions (17% versus 5%), ostensibly testing and probing Jill Watson's capabilities.

We also find that there is a difference in the performance of Jill Watson on the two platforms. The accuracy, i.e., the fraction of questions that Jill Watson answered correctly is 82% on Edstem and 68% on Canvas LTI. This is probably due to students asking better-formed questions on the public forum as they learn vicariously from questions posed previously by other students.

Going forward, we plan to probe further the student engagement with Jill Watson as well as improve its accuracy to enhance its utility in the online classroom.

Click to watch demo of Jill Watson

SAMI

Learning is not only a cognitive process but also a social and emotional process. Many online students often feel socially (and thereby also emotionally) isolated because of the distributed asynchronous nature of most online classes. SAMI (**Social Agent Mediated Interactions**) is a virtual agent developed at the **Design & Intelligence Lab** led by **Professor Ashok Goel** at Georgia Tech. It uses AI and natural language processing to enhance social connectedness in online classes: based on students' self-introductions posted in a class discussion forum, it connects them based on shared locations, interests, and academic goals.

SAMI is currently deployed in five computer science classes with about 2,500 students at the Georgia Institute of Technology. While SAMI is primarily deployed in class discussion forums, we are currently also testing a Slack-based deployment to explore its effectiveness in a more conversational and informal format. Our primary research questions are: Does SAMI help online learners form social connections? If yes, how effective is it in performing this task?

To this end, we have conducted student surveys and gleaned the following insights:

- Students report using shared identity highlighted by SAMI to form connections.
- Highlighting shared identity also helps foster a sense of belongingness, i.e., the feeling that there are others like me in the class.
- Students report mixed opinions on the effectiveness of SAMI in fostering communities:
  - 8% of students in the survey connected with more than 5 students.
  - 20% connected with less than 5 students.
  - 53% did not connect at all.
- Most of our students have familial and work obligations. Along with their academic load, these impact their desire to form social connections. In general, students cite help with academics (e.g., forming study groups and teams for group projects), professional networking, and emotional support (via shared identity and belongingness) as the three common motivators for forming connections.

In another survey conducted in a class with a group project and one that had SAMI deployed, students were asked about how they formed teams:

- 20% of 61 respondents formed teams based on in-person connections; 20% used SAMI's recommendations; 60% used a combination of in-person connections and SAMI's recommendations.
- 70% reached out to connections recommended by SAMI
- 66% responded to a student who reached out to them following SAMI's recommendations.

While it seems SAMI fostered social connectedness and helped students form teams for the task on hand, the relative quality of these teams (SAMI versus non-SAMI based) will be the subject of a future survey close to the end of the semester.

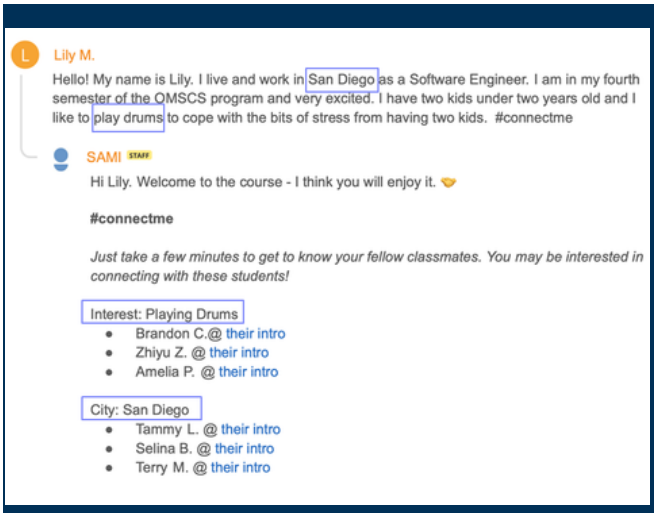

A Screenshot of SAMI's Interactions with a Student

Click to watch demo of SAMI

VERA

Learning sciences often requires conducting experiments in a physical laboratory to practice lab skills. For online learners or commuter learners who have responsibilities that complicate scheduling in the on-campus lab, how could we enable them to participate in lab-based learning activities without the overhead and constraints of a physical lab? **VERA (Virtual Experimental Research Assistant)** is a virtual laboratory for ecology developed at the **Design & Intelligence Lab** led by **Professor Ashok Goel** at Georgia Tech in collaboration with the Smithsonian Institution's Encyclopedia of Life Project to support inquiry-based learning. Students can use VERA to generate hypotheses for explaining given data, construct conceptual models, simulate the models to evaluate them, and validate the simulation results against data.

VERA Research Mode

We recently developed software that allows researchers to conduct A/B experiments in VERA. Researchers can turn on or off certain VERA features for diverse groups to test the following conditions:

Advanced Parameters - During the agent-based simulation, there are parameters available to users to customize a "biotic" component.
Cloning - One way a user can use VERA is by creating simulations based on examples. The cloning feature enables users to "make a copy" of an existing exemplar and test different results of using the software with or without the ability to create a project based on "Exemplar Models".
Exemplar Models - The exemplar menu option allows users to see examples that VERA administrators have created. This option is different from cloning picture because it completely removes the "Exemplars" option from the menu and makes it so that users cannot see the examples.
Lookup EOL - EOL stands for the Encyclopedia of Life at the National Museum of Natural History. The VERA team has partnered with EOL to allow access to their database and extract information on organisms to populate biotic parameters automatically. This feature allows a researcher to test differences in user success based on having to choose their parameters or having them filled automatically.
Simulation - The VERA program enables a user to create a concept model in the model editor, have a simulation run automatically, and then report the results of component numbers in a graph.

VERA in NGTC Classes

AI ALOE formed a partnership with North Georgia Technical College (NGTC) in 2022 and deployed VERA in two courses, ESCI 2030: Forest, Stream and Wetland and ESCI 2060: Advanced Wildlife Management. VERA was deployed in ESCI 2030 in the summer 2022 semester, where thirteen enrolled students were (unexpectedly) given a modeling task of a fish ecosystem native to North Georgia using VERA on the final exam. Students reviewed model templates developed by VERA team members prior to developing their own. Specifically, students were instructed to develop the model step-by-step without using a template model and perform simulations throughout. Altogether, the thirteen students made 25 models.

An early finding shows that students spent on average over 11 minutes on the models and developed models that were 2.5 times more complex (as measured by the total number of components and relationships among components on the model canvas) than needed to solve the question. Based on the initial findings, several potential research opportunities are presented, including how students use VERA on exams vs. a more relaxed setting.

NGTC will continue to utilize VERA in courses in the summer and fall of 2023 with the new Teacher Mode functionality that will allow course instructors to manage course rosters and view analytics on students in a web dashboard. Metrics such as time spent and complexity of models built will be available at the individual and aggregate levels. While the new Teacher Mode functionality will bolster instructors' abilities to implement VERA effectively, it will also allow the automated collection of data that can be used to analyze learning.

Click to watch a demo of VERA

iTELL

iTELL (**Intelligent Textbooks for Enhanced Lifelong Learning)** is under development at the **Language and Educational Analytics Research Lab (LEAR Lab)** led by **Dr. Scott Crossley** at Vanderbilt University. The goal of the iTell framework is to create a digital textbook that will provide students with opportunities to interact with learning materials and have AI systems that generate feedback for students specific to their learning. The iTELL framework currently integrates traditional texts, hyperlinks, videos, and quizzes to promote learning. As a learning analytic system, iTell captures events and text input from learners, and stores learner data in a database. The events captured include click-stream data, key-stroke logging, temporal data, and reading focus data. Unique to the iTell framework is the inclusion of interactive summarization sections that promote active reading enhanced through AI feedback.

iTELL

In the last quarter, we have developed an automatic feedback system for student summaries that are required at the end of each section of a textbook. Prior to moving to the next section, students need to successfully summarize the current section of the textbook. Success is based on natural language processing models that assess whether the summary produced by the student accurately captures the content of the section and contains appropriate wording and whether the students produce sufficient keywords to indicate understanding. The models will provide formative feedback to the students personalized to their summaries. This feedback includes how well they captured the gist and the main details of the section, their appropriate use of objective language and ideas from the section, and key ideas that the students may have missed. The feedback will help scaffold student learning through the explanation of assigned material to ensure greater text comprehension.

We are working with partners at Georgia Tech and TCSG to assess the feedback mechanisms integrated into iTell in terms of useability and effectiveness. We are also developing increased capacity within the system to automatically generate questions and answers to those questions based on individual learner reading patterns.

Summarization Models

We have developed numerous models to automatically score text summarizations within iTELL and within the SMART tool. Our first steps in assessing text summarization involve the use of junk filters to ensure that writers are effortful. These filters start by ensuring that the summary is between 50 and 250 words in length, is written in English, is on topic, and is not plagiarized. To develop appropriate models, we have been pursuing several avenues of inquiry. First, we have fine-tuned large language models that can compare the text of the summary with a keyword analysis of the section. The keyword analysis will help us automatically highlight sections of the text that users need to refer to ensure they are capturing important content within the section. Second, we have developed automatic summary scoring systems by training transformer neural network models on previously graded summaries to predict the score of a student summary. Here, the summary, along with the textbook section summarized, is fed into a large language model (LLM) developed specifically to score summaries on two dimensions – content and wording. The LLM is based on a Longformer model that is fine-tuned on over 5,000 hand-scored summaries. Results from testing the model on summaries that were withdrawn from the training set are promising, and early post-hoc analysis indicates that the LLM can accurately provide feedback on summary content and wording on non-trained topics.

Our next step is to test the model accuracy on textbook section summaries written by > 100 writers recruited on Prolific by comparing human-assigned scores to the model predictions. Our final objective for year 2 is to work with the SMART team to integrate the summarization algorithms above into SMART so that it can provide feedback to learners about the quality of summarization produced within the technologies.

Participatory Design for Human Well-Being

Within AI-ALOE, the team **Participatory Design for Human Well-being** led by **Dr. Michael Hoffmann** ensures that all AI technologies developed by the ALOE Institute will be designed in a way that takes the well-being of users, unintended users, and stakeholders into account. Through user and stakeholder engagement, the Participatory Design Team gives those a voice in the design process who will be affected by the AI-ALOE technologies. The participatory design team follows the IEEE Standard "IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being' (2020)." During this quarter, it has developed a comprehensive roadmap, conducted and analyzed a focus group meeting, and created a work plan for publications and presentations.

Participatory design in general is a method of designing products, services, and systems that involves users and stakeholders in the design process and participatory design for human well-being focuses, more specifically, on effects on well-being. The goal is to create successful designs that meet the expectations, needs, and preferences of the people who will be using them and those who are affected by them. The participatory design method typically involves a collaborative and iterative process that includes multiple stakeholders, such as designers, users, experts, and community members in the design process. There are many ways to conduct this process, including workshops, focus groups, interviews, surveys, and other methods that allow participants to directly express their opinions and provide direct feedback. The benefit of participatory design for human well-being is that this method not only leads to greater user engagement and ownership, but also increases social inclusion, equity, and accessibility as well as the early identification of possible negative effects on well-being.

1. Developed a comprehensive roadmap: Following the "IEEE Recommended Practice," we developed a comprehensive roadmap below that visualizes definitions, goals, and all activities step by step. Based on the IEEE standard, eleven main activities will be applied for each technology.
2. Analyzed SMART's focus group meeting: A focus group meeting was held for the SMART technology on November 3, 2022, in cooperation with Columbus Tech, a college within the Technological College System of Georgia (TCSG). The analysis of the video-recorded meeting resulted in a 16-page report.
3. Created a work plan for publications and conference presentations.

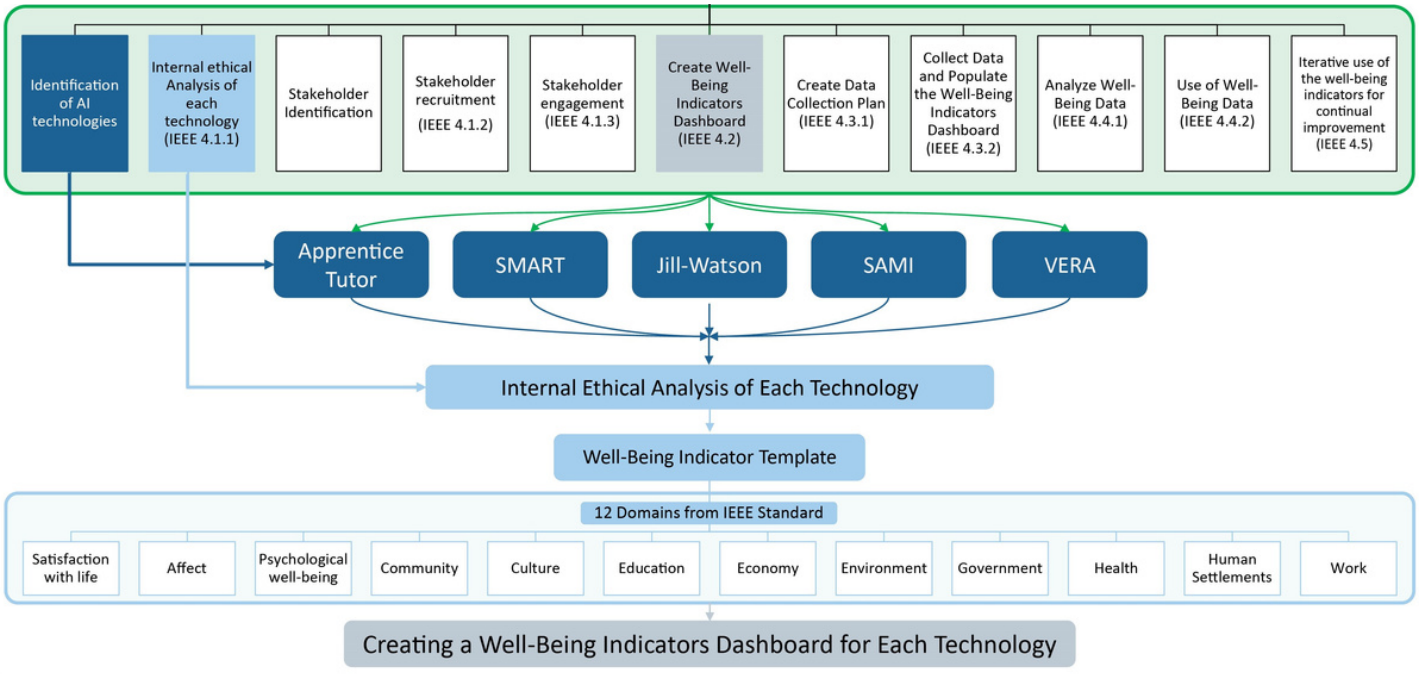

Data Visualization

The **Visual Analytics Lab** at Georgia Tech has engaged in two primary lines of research as part of AI-ALOE. Consisting of **Alex Endert, Adam Coscia, and Grace Guo**, their research has enabled both explorations of how visualizations can be used to help researchers, instructors, and students analyze data collected by intelligent tutors and how the design and development of visualizations can debug and interpret large language models.

Visualizing Intelligent Tutor Data

We built **several visualization prototypes** with classroom data from prior deployments of AI tutors at collaborating institutions. These prototypes help us identify key behavior patterns in how students approach math problems, which in turn reveals insights into problem-solving strategies and challenges. Looking forward, we will work with teachers to better understand how visualizations and data analysis can improve AI tutor use in classrooms. For example, how can teachers be empowered to explore patterns in student behavior through the interactive specification of complex regular expression queries? Our initial visualization prototypes will be refined through a series of participatory design meetings and integrated into the AI-ALOE tutor platforms. By "closing the loop" on the AI tutor development cycle, we aim to provide teachers with tools to analyze AI tutor data and contextualize this analysis with their own knowledge about their classrooms.

Visualizing Large Language Models

As AI-ALOE incorporates large "black-box" language models like BERT in their AI technology, we are also developing novel visualizations that can "open the box" and interpret how and why these language models work. By understanding how and why our models work, we can improve their performance in several ways, such as reducing bias, mitigating stereotypes, learning important domain knowledge, and more. Our latest development in this area is **KnowledgeVIS**, an interactive visualization interface that compares fill-in-the-blank sentences to reveal learned associations in language models.

Consider asking BERT to predict what should be filled in the sentence, "The man worked as a ___." BERT gives responses such as carpenter and miner. However, when asking, "The woman worked as a ___", BERT responds with waitress and prostitute. By intuitively formatting multiple fill-in-the-blank sentences as prompts to the model, we can reveal learned associations that help us interpret how the language model is performing.

The interactive visualizations in KnowledgeVIS help summarize the likelihood and uniqueness of predictions, compare sets of predictions between sentences, and identify patterns in predictions across multiple sentences. KnowledgeVIS currently demonstrates several capabilities, such as uncovering harmful gender biases, adapting prompts to medical domain knowledge, and even supporting open-ended prompts with subjective answers. With new techniques for qualitatively evaluating language model performance, KnowledgeVIS contributes to important machine learning interpretability research at AI-ALOE.

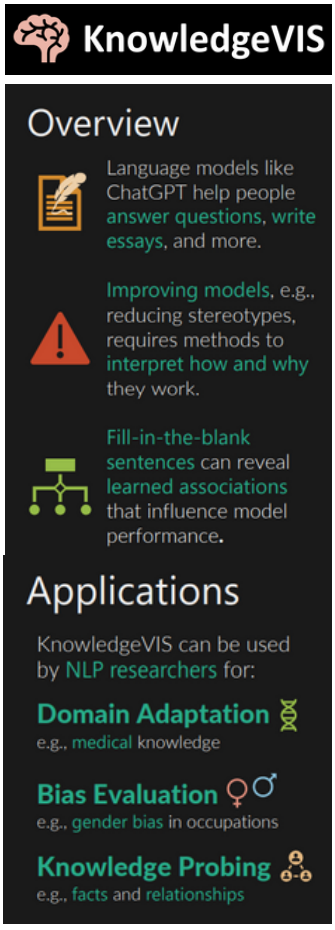

Theory of Mind

Theory of mind is a cognitive and social construct that refers to a human's mental model of another human's mind and mental state. ToM is generally considered to be the basis of human communication. Humans and AI agents, however, do not have a mutual theory of mind (MToM): most humans don't understand much of how AI agents work or think, and, similarly, AI agent's capacity to understand human behavior too is limited. **Professor Ashok Goel** and his **Design & Intelligence Lab** constructed the overall framework of Mutual Theory of Mind (MToM, consisting of three states and three elements. The three elements are constantly shaping the mutual understanding between humans and AI during three stages of the human communication process (see the framework below).

A Framework of Mutual Theory of Mind

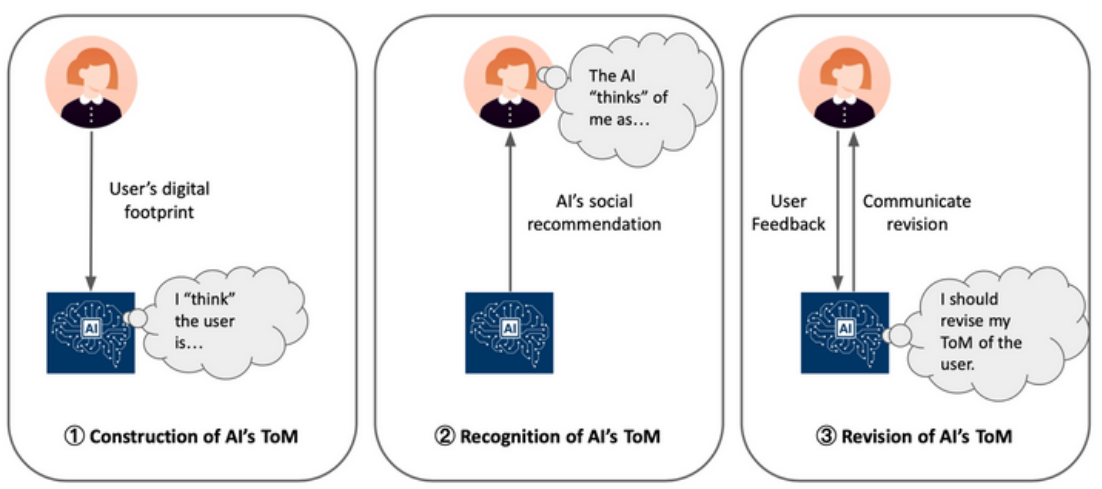

Guided by the MToM framework, we continue to conduct the following three studies to further understand and explore its potentials:

1. **MToM Literature Review**: We are conducting an ongoing literature review to understand how "Theory of Mind," a concept of social science roots, has been applied and leveraged in computer science literature. This will help us better position and situate the value and contribution of the MToM framework. Through a scoping review of the existing CS literature and social science literature, we aim to address three research questions: (1) What is the role of ToM in human-human social interaction according to social science literature? (2) How is ToM defined, applied, and addressed in human-AI interaction according to CS literature? (3) How can CS researchers better utilize the concept of ToM to enhance human-AI social interaction?
2. **Empirical Study on Recognition of AI's ToM**: This study takes place in the context of using SAMI for team-matching purposes and aims at understanding three research questions: (1) What do students think of personality-based algorithm-facilitated team matching? (2) How do students make sense of AI's understanding of them in algorithm-facilitated team matching? (3) How would inaccurate AI inferences impact students' perceptions of AI in algorithm-facilitated team matching? As of Spring 2023, we finished the initial pilot study and the study preparation for the recognition study, and we have begun the data collection for the recognition study. So far, we have conducted study sessions with nine participants.
3. **Development of MCM for Revision of AI's ToM**: We are also making progress in building and developing the metacognitive module for SAMI to understand and explore the final stage of the MToM framework--- the revision of AI's ToM. Modeling AI's metacognition process is critical to increase the transparency of AI's revision process and improve mutual understanding between humans and AIs. During this development process, we are also incorporating the TMK (Task-Methods-Knowledge) model to help model SAMI's cognitive process and metacognitive process.

Team Member Spotlight

Chris MacLellan

Get to know **Chris MacLellan**, an assistant professor in the School of Interactive Computing at Georgia Tech where he runs **Teachable AI Lab (TAIL)** His lab aims to improve the understanding of how people teach and learn and to build AI systems that can teach and learn like people do. His recent work explores the development of computational models of human learning and how these models can support the development of effective learning technologies at scale.

Chris received his PhD in Human-Computer Interaction from the Human Computer Interaction Institute at Carnegie Mellon University. Prior to his work at Georgia Tech, he was an assistant professor at Drexel University in the Information Science and Computer Science Departments. He received the Faculty Research Excellence Award presented by Drexel University's College of Computing Faculty Research, which recognizes recipients for their scholarly research and prolific academic and professional contributions to the University.

Revolutionizing AI in Education with Teachable AI

By Chris MacLellan

We are working towards creating a teachable AI technology that lets teachers build and modify intelligent tutoring systems to support their classes. The aim is to build an interactive AI agent that lets teachers author tutors through teaching rather than programming. To create a tutor for a new topic, the teacher provides the agent with examples of how to solve problems and corrects it problem solving on novel problems—similar to how a teacher would teach a human student. Once the agent learns how to do the task, it can then provide one-on-one instruction and coaching to students at scale.

We are currently building a platform called Apprentice Tutors that let teachers create their own tutors using this technology. We are currently focusing on deploying tutors into classrooms through Blackboard, and actively testing tutors to support college algebra classes at TCSG. This past year, more than 2000 TCSG students have had access to tutors through our platform.

While we are working on the foundational AI research about how to teach machines, we are also collaborating with teachers to build effective tutors. As students interact with the tutors, we use their data to assess their knowledge and to personalize tutor behavior to their current learning state. We also use this data to analyze class performance and assess student learning.

One of the challenges we face is building technologies that are accessible and personalized to teachers and students. We are working with teachers to build tools they can use to create personalized educational technologies for their classes. On the technical side, we're trying to advance AI technologies to enable learning from teachers with just a few examples. Our goal is to teach agents with as few as 10 demonstrations so that they can teach students effectively. This is a foundational AI challenge that we are working to overcome.

Meet Our Students

Langdon Holmes

Meet **Langdon Holmes**, an AI-ALOE member and a second year Ph.D. student in the Department of Psychology and Human Development at Vanderbilt University where he studies language learning and educational technology. He is advised by Dr. Scott Crossley in the **Language and Educational Analytics Research Lab (LEAR Lab)**. Their work is focused on using data science to generate insights about language, learning and education that support AI-driven theory, applications, research and interventions.

Langdon's research areas are learning analytics, second language acquisition, transformers. He was awarded a $30k grant from the Learning Agency to create a dataset of student-generated text labeled for personally identifiable information. The project will help address private information appearing in learner data, which helps prevent this type of data from being shared between researchers.

What are your main research areas at AI-ALOE?

My two main research areas are “deidentification of student data for open science” and “learning analytics with natural language.”

What drives your research with AI-ALOE?

At ALOE, we believe that AI has the potential to improve learning outcomes for adult learners. Adult learning will allow individuals to reskill in response to changes in technology and other factors that affect their employability. Adult learning also may uncover hidden talents in individuals who did not meet their academic potential as adolescents. Ultimately, I am driven by the idea that learning will become a bigger part of people’s lives, giving everyone more opportunities to develop valuable skills and reach their fullest potential.

How did you develop an interest in your research?

As an undergrad, I studied comparative literature with a particular interest in Latin American poetry. I thought I would become an English teacher and spent over five years teaching English. During that time, I became increasingly interested in linguistics because it allowed me to give my students better feedback. When the pandemic struck, I was nearing the end of my Master’s degree in applied linguistics, and I found myself with plenty of spare time to learn how to write computer code. I became fascinated with large language models and began learning about them as much as possible. I still have a deep interest in education (and especially language learning), but now I am more interested in building tools for learners and educators. As a teacher, I could only help 30 students at a time (on a good day), but with educational technology, there is potential to help transform education at a societal level.

What are your hobbies and some fun facts about you?

I like board games, video games, and dinner parties. Here's a fun fact: I sold my car and bought a one-way plane ticket to Chile when I graduated college. I stayed there for two years (eventually finding work as an English teacher) before applying to graduate school.

Outreach & Collaboration

AI in Education Symposium: AI-ALOE's One Year and Beyond

AI-ALOE celebrated its one-year's anniversary by hosting a symposium on Nov. 11, 2022. Eighty-two researchers, educators, and students joined us to discuss the past, present, and future impacts of using AI in education.

AI-ALOE Leadership Team Participated in SAIL

AI-ALOE leadership team attended the Summit for AI Institute Leadership (SAIL)from Dec. 13-15, 2022. AI-ALOE's Executive Director Ashok Goel moderated and participated in a panel discussion on expanding AI in education and broadening participation during the summit.

AI-ALOE at the Engage AI Institute Forum

On Dec. 16, 2022, AI-ALOE participated in the Engage AI Institute Forum held at the Computer History Museum in Mountain View, CA. It was a one-day gathering of the NSF-funded AI Institutes addressing education and people near Silicon Valley with expertise and interests in AI in Education.

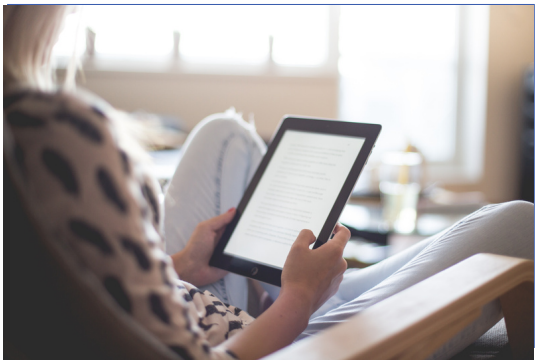

AI-ALOE Conducted a Large-Scale Digital Textbooks Survey

AI-ALOE partnered with Georgia Tech, Georgia Highlands College, and the Technical College of System of Georgia to conduct a large-scale survey on how students use digital textbooks and what features they wish to have to help them learn beyond reading text. Dr. Chaohua Ou led the project, and more than 3,700 students completed the survey. Their responses provide insight into how we can design and develop intelligent learning content.

AI-ALOE to Lead an AERA Panel

Dr. Chris Dede will moderate a panel at the Annual Conference of the American Educational Research Association(AERA) on April 15, 2023, in Chicago.The panel will bring together the three AI Institutes on learning and education. Dr. Ashok Goel will be a panelist, and Dr. Amy Baylor from the National Science Foundation will be the discussant.

Publications & Presentations

Publications

An, S., Hammock, J, Rugaber, S., and Goel, A. (2022) Contextualized Access to Large-Scale Domain Knowledge for Conceptual Modeling of Agent-Based Systems. In *Proceedings of the 10th Annual Conference on Advances in Cognitive Systems*, pp. 578-599.

Dede, C. (2022). The Coming Sea-Change in Teacher Education. *Journal of Technology and Teacher Education*, *30*(2), 117-125. Waynesville, NC USA: Society for Information Technology & Teacher Education. https://www.learntechlib.org/primary/p/221170/.

MacLellan, C.J., Stowers, K., Brady, L. (2022). Evaluating Alternative Training Interventions Using Personalized Computational Models of Learning. *Advances in Cognitive Systems*, *10*, 1-18.

MacLellan, C.J., Matsakis, P., & Langley, P. (2022). Efficient Induction of Language Models via Probabilistic Concept Formation. In *Proceedings of the 10th Annual Conference on Advances in Cognitive Systems*.

Morris, W., Crossley, S., Holmes, L., Ou, C, Dascalu, M., & McNamara, D. (under review) Formative Feedback in Intelligent Textbooks using Large Language Models. *Journal of Artificial Intelligence in Education.*

Morris, W., Crossley, S., Holmes, L., & Turmbore, A(2023). Using transformer language models to validate peer-assigned essay scores in Massive Open Online Courses (MOOCs). in Proceedings of the 13th International Learning Analytics and Knowledge Conference (LAK2023). pp. 315-323. doi.org/10.1145/3576050.3576098

Radu, I., Dede, C., Wang, J., Nie, G., Bhola, K., & Scuzzarella, M. (2022). Using virtual environments to reveal teacher bias towards students' socioeconomic status. *2022 8th International Conference of the Immersive Learning Research Network (iLRN)*, pp. 1-8. Doi.org/10.23919/iLRN55037.2022.981595

See a full list of publications and presentations on our website

Presentations

Chris Dede gave a keynote address about AI-ALOE for the AI and Edu workshop sponsored by the PALS group at the University of Massachusetts Amherst in December 2022.

Chris Dede discussed on the Harvard EdCast how artificial intelligence (AI) has the power to address some of the biggest challenges in education in January 2023.

Ashok Goel moderated and participated in a parallel session on expanding education and broadening participation during the Summit for AI Institutes Leadership in December 2022.

At the Future of AI & Data-Enabled Discovery Workshop hosted by IDEaS at Georgia Tech in February 2023, Ashok Goel outlined the Institute's goal to build innovative technologies advancing adult online education for lifelong learning and contributions towards shaping the future of the workforce.

Christopher MacLellan visited UC-Boulder to give a talk in November 2022 on Teachable AI: A cognitively inspired and human-centered approach to the knowledge transfer problem.

Christopher MacLellan gave a virtual talk, *Efficient Induction of Language Models Via Probabilistic Concept Formation,* during the tenth annual Conference on Advances in Cognitive Systems (ACS) held in Arlington, Virginia, November 19-22.

Min Kyu Kim and his team have three papers accepted for presentation at the 17th International Conference of the Learning Sciences/Computer-Supported Collaborative Learning (ICLS/CSCL-2023) in Montréal, Canada.