



THE AI-ALOE SPOTLIGHT



Ashok Goel
Director

We celebrated the two-year anniversary of the AI-ALOE Institute on November 1st. Leading up to this significant milestone, we welcomed a new assistant to the executive and a program support coordinator. Over the past year, our research institute accomplished several noteworthy mile-stones, including the completion of our annual report and annual review by the National Science Foundation (NSF). Now, as we embark on year three, our institute is actively engaged in more than 14 research projects that encompass a wide range of areas and education technologies.

Our mission to innovate AI-based transformative models for online adult learning is making significant progress in both our Use-Inspired AI and Foundational Research into AI research thrusts, as well as Research Design.

In this edition of our monthly newsletter, which serves as a consistent source of information on the institute's growth and achievements, we introduce you to one of our exceptional students, Chelsea Wang, who is actively involved in the Design & Intelligence Lab (DILab). Chelsea's work centers around the Mutual Theory of Mind Framework in Human-AI Interaction. Additionally, we showcase the research efforts of our Assistant Director for Outreach, Scott Crossley, who is focused on utilizing data science and AI to explore cognition, particularly in the field of natural language processing.

This issue also highlights our outreach activities and engagement in various conferences and workshops over the past few months, where our members actively participated and contributed.

As we look ahead to the future, we are eager and enthusiastic. As we proudly enter year three, I would like to express my gratitude to Georgia Tech and the NSF for their continued support. I would also like to extend my sincere appreciation to all the members of AI-ALOE for making this possible.

FALL 2023 NEWSLETTER

Table of Contents

Page 1
Director's Message

Page 2
A4L

Page 3
Apprentice Tutor Builder

Page 4
Conversational Courseware

Page 5
ITELL

Page 6
Interactive Videos (IVY)

Page 7
Participatory Design

Page 8
SAMI

Page 9
SMART

Page 10
VERA

Page 11
Visualization

Pages 12 -13
AI-ALOE News & Events

Page 14
AI-ALOE AT SAIL 2023

Page 15
Collaboration & Outreach

Page 16
Team Member Spotlight

Page 17
Student Highlight

Pages 18 - 19
Publications and Presentations

AI-ALOE Social Media



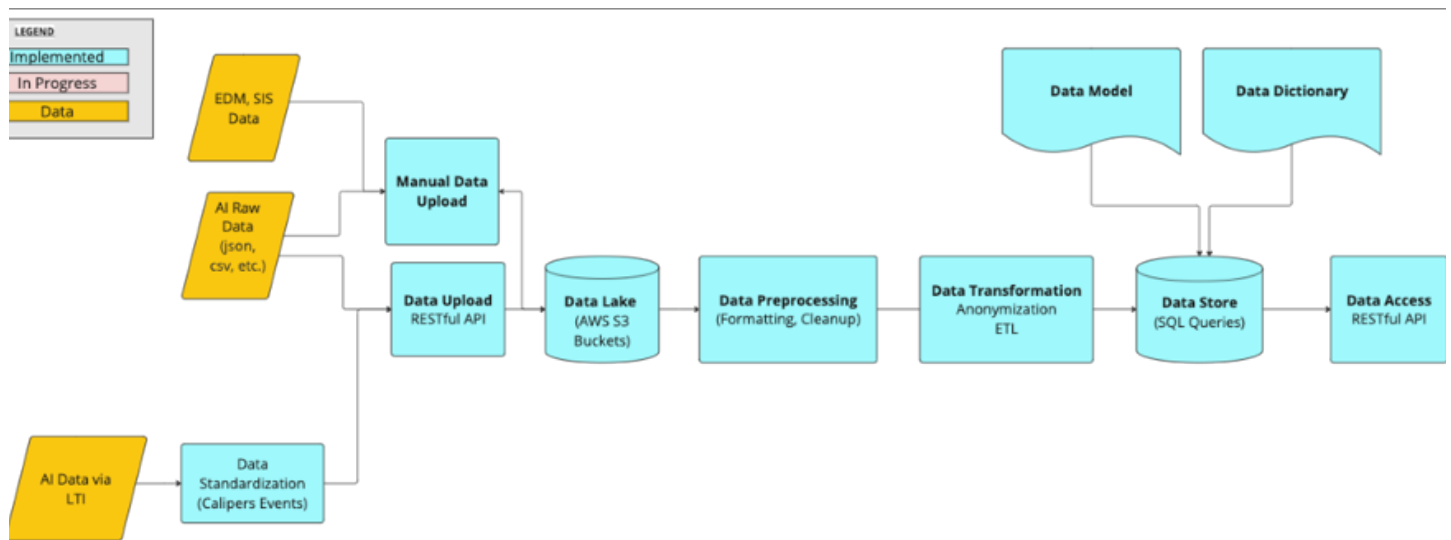
A4L Technology Infrastructure

The A4L (Architecture for AI Augmented Adult Learning) Data Engine optimizes data management with features like accommodating various data formats, automated or manual data ingestion, ETL processes for data quality, a centralized data repository, and robust data access. It seamlessly integrates data from diverse sources to support micro-, meso-, and macro-learning, and serves various applications including AI/ML agent personalization, data analytics, and data visualization.

The products developed by the A4L team during this quarter include:

- **A4L Data Engine:** Our flagship product, the A4L data engine, remains a versatile tool for AI/ML agent personalization, data analytics, and data visualization.
- **Data Model Summary:** A preliminary version of the conceptual data model has been created for three distinct technologies: JW, SAMI, and VERA.
- **Data Integration:** The platform seamlessly combines data from different sources, supporting micro, meso, and macro learning.
- **Data Collection:** We now collect and manage various data types, including learner performance, AI interactions, survey responses, demographic information, grades, and social connections.
- **Access and Sharing:** We've developed access and sharing agreements, data retention criteria, security measures, and compliance procedures to ensure data integrity.
- **Governance and Documentation:** Our commitment to governance, compliance, and meticulous documentation provides transparency into data management.

A4L: Conceptual Architecture



Apprentice Tutor Builder

Towards the goal of enabling teachers to create and personalize tutors easily without the need for AI expertise, we have created a new technology called Apprentice Tutor Builder (ATB). This system currently consists of two parts: an interface builder (see figure 1 below), which lets teachers create scaffolded tutor interfaces using a drag and drop approach, and a teachable AI agent (see figure 2 below), which lets teachers create a personalized AI expert model to power the tutor by teaching via demonstrations and feedback directly in the tutor interface they created (similar to how they would teach a human student in the same interface). We also have a line of research (Verbal Apprentice Learner or VAL) that will let teachers also use language instruction in addition to demonstrations and feedback when teaching the agent.

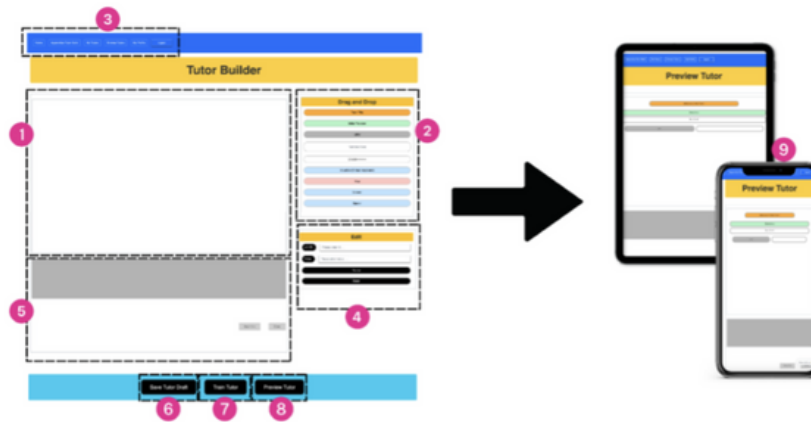


Fig. 1. An example of the Apprentice Tutor Builder (left) and built tutor interfaces (right). Left (1) drag and drop pane (2) draggable objects and components (3) header bar to navigate to other tutor builder screens such as profile, tutors, tutor marketplace (4) edit component pane to change an objects properties (5) hint box and corresponding buttons (6) save button (7) train button (8) preview tutor button. Right. (9) previewing built tutor on different mobile devices including an iPhone and iPad.

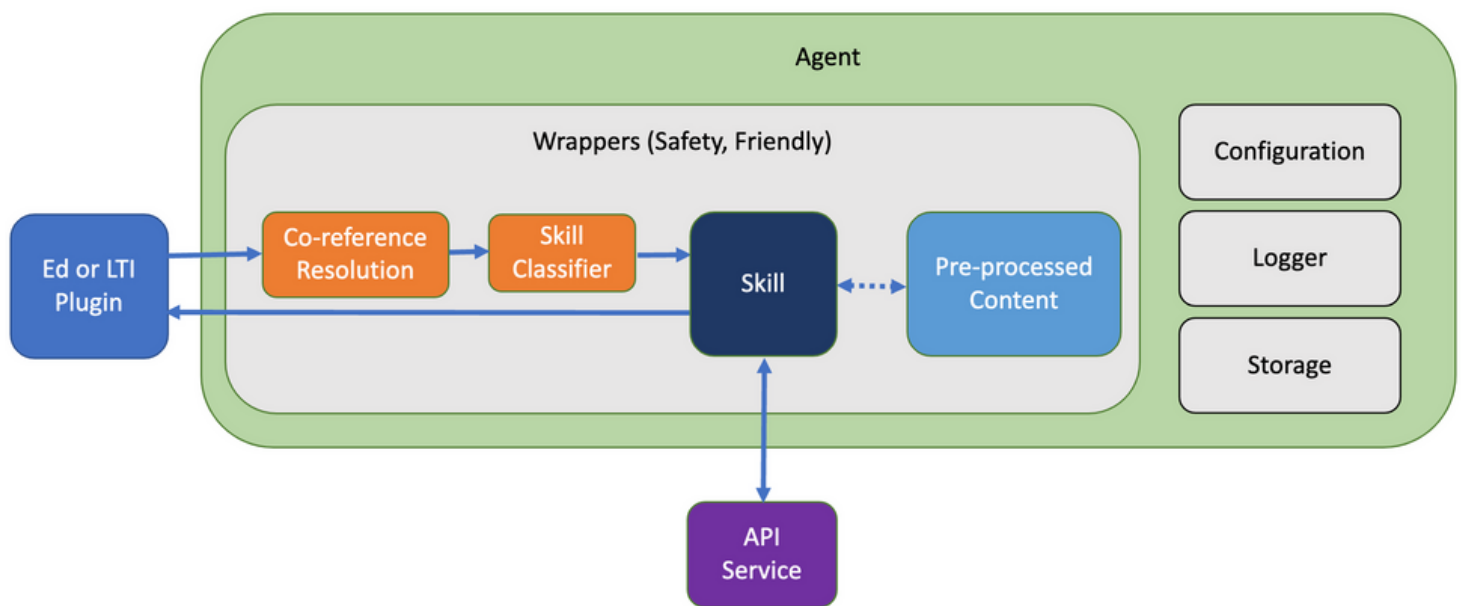
Figure 1. An example of the Interface Builder (left) and constructed tutor interfaces (right). Left (1) drag and drop pane (2) draggable objects and components (3) header bar to navigate to other tutor builder screens such as profile, tutors, tutor marketplace (4) edit component pane to change an objects properties (5) hint box and corresponding buttons (6) save button (7) train button (8) preview tutor button. Right. (9) previewing built tutor on different mobile devices including an iPhone and iPad.



Figure 2. The teachable AI agent training interface for Apprentice Tutor Builder
Over the summer we performed a user study with 14 participants (13 teachers, 1 non-teacher) to evaluate the effectiveness and usability of ATB. The study consisted of having these users create two tutoring systems. For each tutor they needed to create a tutor interface in the interface builder and author an expert model for the tutor by teaching an AI agent via a series of multi-modal interactions. Quantitative measurements were collected via user logs and qualitative measurements were collected via a semi-structured interview. We derived metrics from the quantitative data to assess the efficiency and effectiveness of the ATB technology. Additionally, we analyzed the qualitative data using affinity diagramming to generate insights regarding users' experience and reaction to our technology. From this study, we were able to determine critical design features required for ATB for effective use which translated into design recommendations for systems such as ATB. We currently have a paper under review at ACM CHI 2024 detailing this study.

Conversational Courseware

Jill Watson is intelligent courseware that engages students in conversations with textbooks, video transcripts, presentation slides, class syllabi, and other educational materials. It thus enhances both cognitive engagement (through sustained conversations with educational materials including textbooks) and teaching presence (by relating different educational materials to one another and thus making them more meaningful to the students). It uses LLM-based generative AI (ChatGPT) to answer student questions regarding instructor-approved courseware anywhere and anytime. Jill is designed to mitigate the problem of hallucinations commonly seen in LLM output by grounding its answers in the information extracted only from documents in its knowledge base versus much wider, often erroneous or irrelevant, knowledge that is parameterized in LLMs from their training dataset. Jill also has guardrails against undesirable and hateful speech, as well as questions unconnected to the course content.



In Q4, Jill achieved several milestones:

1. At the end of August 2023, we introduced Jill to an online English composition class of about 125 students at Wiregrass College, a 2-year technical college in the Technical College System of Georgia (TCSG) that accepts all students who apply. This is the first time Jill has been deployed outside of the Georgia Institute of Technology. In contrast, Jill is also deployed in an online class on AI with about 500 students at the very selective Georgia Tech Online Master of Computer Science (OMSCS) program. The two courses present very different demographics, and student readiness, motivation and achievement levels, which will allow us to gather data to explore Jill's impact on learning in diverse student bodies.
2. Jill can now engage in conversations based on any textbook, user guides and manuals, presentation slides, and other natural language documents provided by the instructor. For example, in the English composition class at Wiregrass College, Jill is answering questions regarding the Purdue MLA writing and style guide and a 250-page open educational resources (OER) textbook "Successful College Composition." In the AI class at Georgia Tech, Jill is answering questions related to a 357-page e-book. Previously, Jill's ability to answer was confined to questions related to course syllabi and schedules.
3. Jill is now fully conversational. Students can have a one-on-one conversation with Jill thus making their experience more engaging and useful.

As AI transforms how people work and the skills needed to accomplish this work, intelligent texts that integrate AI to make them more engaging and effective will become important tools for learning. K-12 and tertiary education are expanding into effective lifelong learning systems that follow learners as they reskill and upskill throughout their careers. Intelligent texts will play a key role in making education and training more accessible, affordable, efficient, and adaptable to individual learning needs throughout the process.

The intelligent Texts for Enhanced Lifelong Learning (iTELL) framework addresses this need with a strong focus on cognitive theories of text comprehension. Current AI interactivity in iTELL is based on student summarization and short question generation and answering. Read-to-write tasks like summaries and short answer-questions are a common approach for assessing reading comprehension. By integrating student summaries and question generation and answer scoring within iTELL, we hope to support interactive learning for a variety of complex subjects and provide students with opportunities to better comprehend material.

At the end of each section in a text, students are required to write a summary to assess their comprehension of the material. Two summarization scoring large language models (LLM) are currently used to provide feedback on content and wording. An additional LLM extracts keywords from the text and provides direct feedback to learners about keywords they can include in their summaries to improve their writing. The short answer task encourages students to recall information from subsections. During the preprocessing phase, iTELL uses the ChatGPT API to automatically generate short-answer questions for each subsection as well as reference answers to those questions. Student answers are compared to the reference answers using two embedding models finetuned from BLEURT, and predictions from those models are used to provide feedback to the user on reading comprehension.

Constructed Responses

The ultimate result of workers who can focus on their preferences and talents, learn to do their specialized jobs better, and work in larger organizations is that society as a whole can produce and consume far more than if each person tried to produce all of their own goods and services. The division and specialization of labor has been a force against the problem of scarcity.

Question: What are the three reasons Smith offered for how a group of specialized workers can produce more than the same number of workers who try to produce the entire good or service by themselves?

Trade and Markets

Specialization only makes sense, though, if workers can sell the goods they produce for doing their jobs to purchase the other goods and services that they need to do their own specialization-related tasks.

Text Summarization

Write your summary for this section
You can unlock the next section by submitting a good summary of this section.

What makes a good summary

- Be within 50-200 words long
- Be written in English
- Be on topic
- Not be plagiarized

Scoring details

See a professional summary for this section
Number of words: 0
Disable paste

Submit

- ⓘ Before moving onto the next section, you will need to revise the summary you wrote using the feedback provided. Try to include the following keywords: syntax rule, web pages, assembly language
- ▼ Details
- ✔ Wording: You did a good job of paraphrasing words and sentences from the section and using objective language.
- ✘ Content: You need to include more key ideas and details from the section to successfully summarize the content. Consider focusing on the main ideas of the section and providing support for those ideas in your summary.
- ✔ Topic Borrowing: You did a good job of using your own language to describe the main ideas in the section.
- ✔ Topic Similarity: You did a good job of staying on topic and writing about the main ideas of the text.

Interactive Videos (IVY)

In-video tutors were first proposed by David Joyner and Ashok Goel in Fall 2014 when they embedded 100 tutors in online videos as part of the KBAI course in the OMSCS program (see the two figures for examples of in-video tutors). Despite a successful proof of concept, a primary limitation was the ad-hoc approach used to develop the tutors; limited by the Udacity infrastructure and mostly hard-coding tutor responses for each exercise. The current challenge is to systematically formalize at the methodology, architecture and interface or “infrastructure” such that interactive videos can be generalizable and scalable for online courses that have two main characteristics: (1) the primary audience are adult learners, and (2) the course content focuses on skill-based procedural learning, i.e., knowledge of what set of actions to be taken to do a specific task.

The screenshot shows a 2x2 grid with the top row containing '3' and '0', and the bottom row containing '3' and '0'. A green arrow points from this grid to a larger 2x2 grid on the right, which also contains '3' and '0' in both rows. Below the grid is a question in blue text: "What are the next possible states for each of the current states?". At the bottom, there are two buttons: "SUBMIT ANSWER" and "CONTINUE TO ANSWER".

Try again!

It looks like there's a few problems with your answer. Please note of the issues below; note boxes are numbered from top to bottom.

In the third box, no one moved from the right to the left. Remember, the boat can't travel alone!

You've written the same state twice in the first three boxes. There should be three different possible next states.

Remember, in this exercise we're looking for *all* possible next states, not just the legal ones!

In May 2023, we renamed the in-video tutor project as “Ivy” and started working towards conceptually formalizing the above mentioned infrastructure, specifically what each element entails. Given the overarching goal of moving from an ad-hoc to a systematic approach, we are also developing generalized representations of skill-based problems from KBAI as part of an Ivy coach’s “knowledge bank” using the Task, Methods and Knowledge (TMK) modeling language. Using a TMK model for a given problem, we hypothesize that the Ivy coach would be able to do knowledge tracing to determine the set of steps required to answer a question(s) designed by the instructor. Additionally, using the TMK model for a given problem, if an incorrect response was submitted by a learner, the Ivy coach would be able to trace at what point did a learner perform a misstep and provide appropriate feedback.

The screenshot shows a 2x2 grid with the top row containing '3' and '0', and the bottom row containing '2' and '1'. A green arrow points from this grid to a larger 2x2 grid on the right, which also contains '3' and '0' in both rows. Below the grid is a question in blue text: "What are the next possible states for each of the current states?". At the bottom, there are two buttons: "SUBMIT ANSWER" and "CONTINUE TO ANSWER".

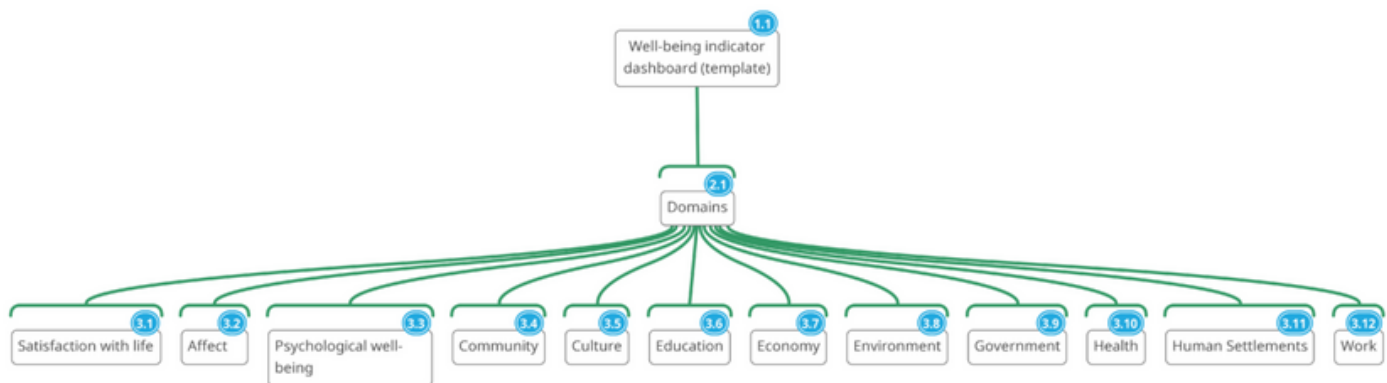
Correct!

Very nicely done! You've written every possible next state. Next, we'll look at how our dumb tester would rule out some of these states.

Participatory Design

Within AI-ALOE, the team "Participatory Design for Human Well-being" ensures that all AI technologies developed by the ALOE Institute will be designed in an ethically responsible manner. The work of the PD team focuses on the issue: How do we realize a responsible design of AI technologies? Our research is driven by the following hypothesis: Realizing the Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being (IEEE Std 7010-2020) ensures the responsible design of AI-ALOE technologies. The PD team uses interviews with AI-ALOE teams that design and deploy AI learning technologies as the method to gather evidence that might support this hypothesis.

A major result of the PD team's work is the gathering of evidence that indicates that this hypothesis can be confirmed regarding the development of SMART technology. In its report (see 3.2 below), the PD team concludes that the stakeholder engagement process that the PD team initiated with the focus group meeting on November 3, 2022, motivated substantial changes in the design of the SMART technology. As Dr. Min Kyu Kim noted in the interview: "most of the changes we have made are related to the comments we have received from the focus group."



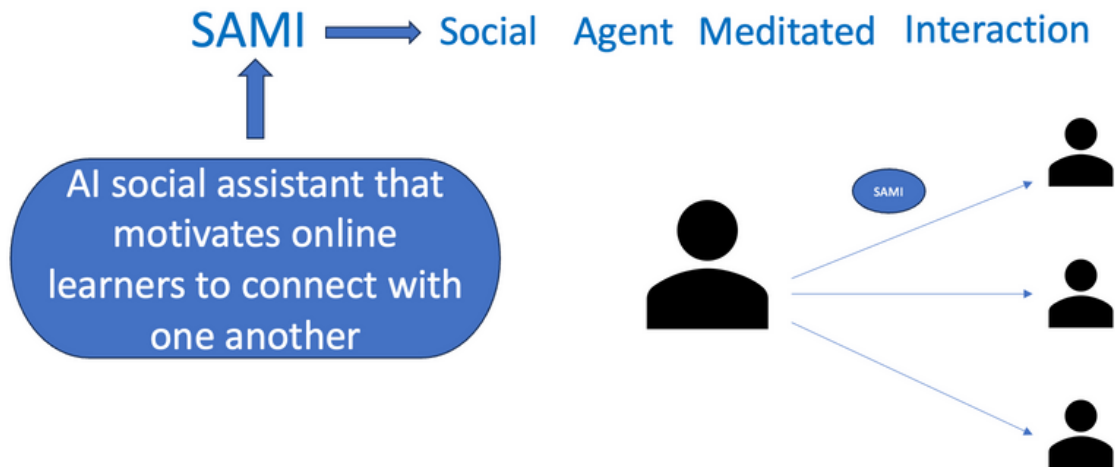
An essential component of the SMART team's approach to address some of the issues that the PD team identified through its stakeholder engagement process concerns the replacement of the current NLP tool (for automatically summarizing texts and creating concept and relation maps) by generative AI. Since this presents a major step in the development of SMART, it should be important to also assess the new version of SMART with a well-being impact analysis.

Some other changes in design of SMART include the following:

- User interface (UI)
- Data visualization
- Providing more guidance for users
- More functions and features for instructors to help them accelerating grading, accessing learners' information

Overall, we see our hypothesis confirmed with regard to SMART: Realizing the Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being (IEEE Std 7010-2020) ensured the responsible design of an AI-ALOE technology.

SAMI



SAMI (Social Agent Mediated Interactions) is an AI-based social agent that aims to enhance social connectedness among students in large online classes which are characterized by geographically dispersed student bodies and asynchronous teaching. The hypothesis underlying this effort is that since learning is as much a social process as it is a cognitive process, increasing students' social presence in the class will result in better learning outcomes. In addition, there may be benefits to student mental health and overall student experience in the class.

SAMI uses generative AI (ChatGPT) to analyze student introduction posts in a class discussion forum such as Ed Discussion or in a conversational forum such as Slack to provide match suggestions based on student location, time zones, hobbies, and academic interest. In Q4:

- We completed a re-write of SAMI to incorporate ChatGPT for named-entity-recognition (NER) to extract student location, time zones, hobbies, and academic interest from students' introduction posts. This approach obviates the need for building and training machine learning models to extract desired entities in the introduction posts, which is a much more complex and time-consuming task. In addition, with some prompt engineering, this list of extracted entities can be easily expanded or modified depending on the student profile in a program, which should serve well to scale SAMI for different target audiences.

SAMI also uses ChatGPT to frame its responses containing match suggestions to students. As a result, SAMI's responses are now more natural, engaging, and diverse as opposed to match suggestions previously embedded in rather stiff hard-coded responses. We expect this to improve student engagement with SAMI.

- We are now able to track if a student acts on SAMI's match suggestions, which should provide us with a direct insight into student engagement with SAMI.
- At the end of August 2023, we deployed SAMI in two AI classes in the OMSCS program. In one class with about 500 students, SAMI is deployed in Slack; in the other class with about 700 students, SAMI is deployed in the Ed Discussion forum. We are collecting data to explore whether a conversational experience afforded by Slack impacts student engagement.

SMART

SMART Experiments in TCSG classes

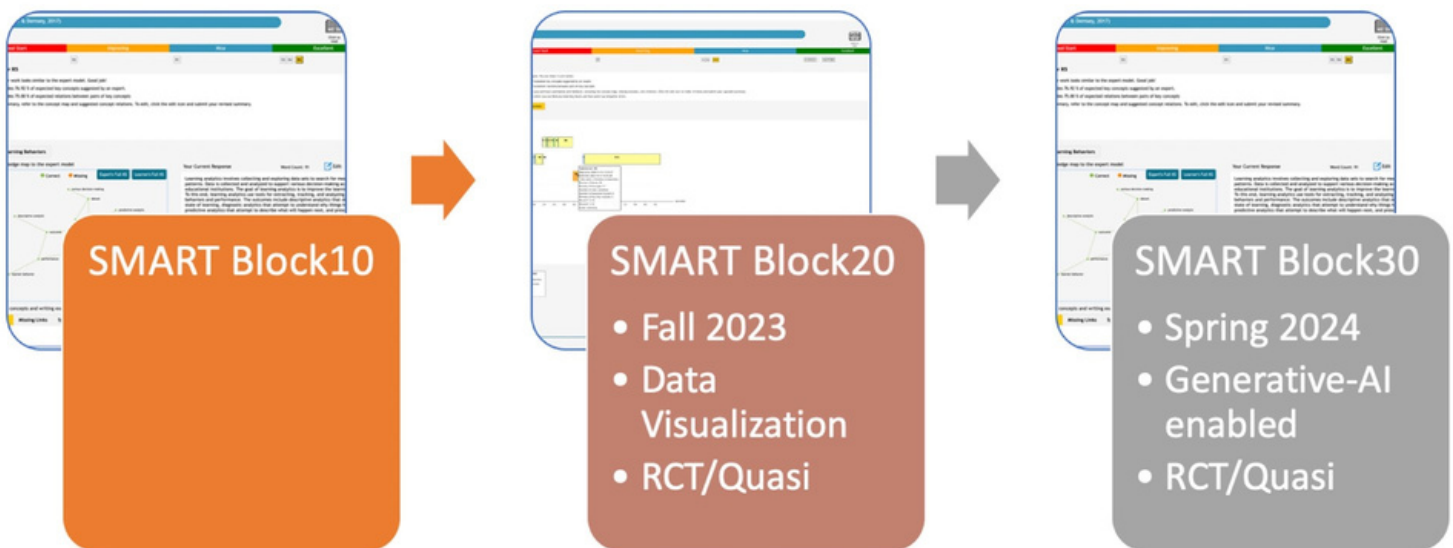
In Fall 2023, SMART (Student Mental Model Analyzer for Research and Teaching) technology is being implemented across four sections of the English course at Columbus Technical College and one section of the Biology course at Chattahoochee Technical College. In the four English classes, we are conducting a quasi-experiment. Two of the classes are using the previous version of SMART (block 10), while the other two classes are engaging with our newly developed SMART (block 20) which is enhanced with data visualization and improved navigation. Specifically, it allows observation of individual learners' behavior patterns regarding their revisions, resulting levels of understanding, and time spent reviewing specific feedback information components. Similarly, in the Biology course, a Randomized Control Trial (RCT) experiment is being conducted, with one set of learners using the previous version of SMART (block 10) and the other using SMART block 20.

SMART's Data Visualization

We have developed a feature to track learners' learning behaviors, integrating a data visualization component into SMART by leveraging D3 JavaScript, a powerful data visualization library. SMART's data visualization feature portrays individual learners' behavior patterns, encompassing their revision history, levels of comprehension achieved, and the time dedicated to reviewing specific feedback components. Additionally, this data visualization functionality not only provides valuable insights into individual learner progress but also allows learners to gauge their performance in comparison to class benchmarks. By spotlighting benchmark data, including essential metrics such as total study duration, revision frequency, and overall achievement levels, it enables learners to become aware of their relative performance levels and motivate them to further engage in the learning activities. This feature fosters self-awareness and encourages a proactive approach to academic progress among learners.

Generative AI-Driven Feedback on SMART

The goal of this work-in-progress project aims to leverage the potential of generative AI to provide richer and more personalized feedback on SMART. In this cycle, we have mainly explored strategies and conducted initial testing for this. Currently, SMART provides feedback to learners by pinpointing missed concepts and relations of their composed summaries. This feedback is derived based on the expert model that the instructor provides. To strengthen the capabilities of SMART, we have designed a framework for a new AI-driven feedback feature using the python library LangChain. LangChain enables large language models to generate context-aware answers and reasoning, and thus, was deemed to fit our needs of our framework.

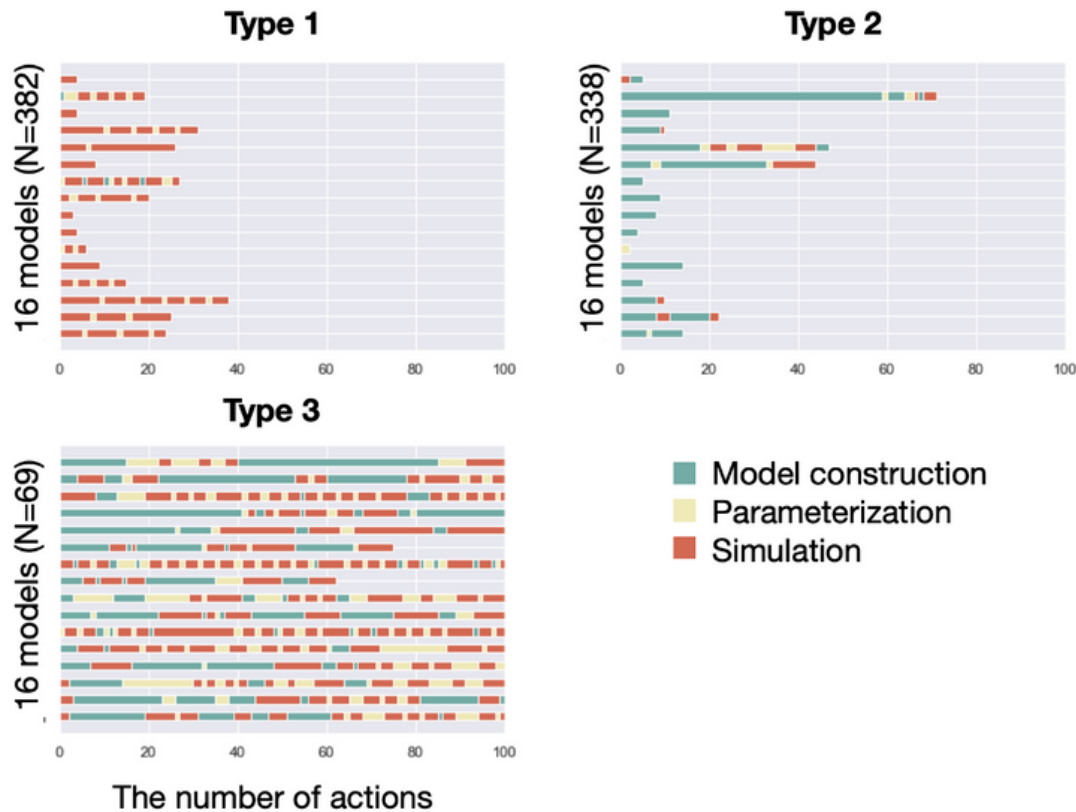


VERA

IVERA (Virtual Experiment Research Assistant) is an interactive environment for inquiry-based learning through conceptual modeling and agent-based simulation. It has been used in a variety of settings such as undergraduate and graduate-level Computer Science, Biology and Ecology classes.

In preparation for Fall 2023, we recently conducted a pilot user study for an initial assessment of our personalized feedback system. Using an interdisciplinary approach, we designed a comprehensive experimental study where first-time users of VERA were asked to narrate their thoughts and actions as they constructed ecological models. Additionally, the participants were interviewed to reflect on their use of/interacting with the coaches. While the overall feedback was positive, the lessons we gathered from the pilot study are not only being used to build more effective coaches and VERA, in general, but they have been instrumental in developing experimental designs for our Fall deployment.

During Fall 2023, we have been refining our personalized feedback system by developing software infrastructure with the goal of assessing the utility and efficacy of personalized feedback to learners solving ill-defined, open-ended problems. In Fall 2023, we launched both AI coaches on the live VERA site. Currently, we are conducting a study to assess the personalized version of VERA in a live classroom (N > 80 students) at Georgia Tech. Using this infrastructure, we will be able to measure how student success compares with and without the help of VERA's feedback. If we are successful, we would be able to demonstrate that a personalized feedback system can be designed for ill-defined contexts.



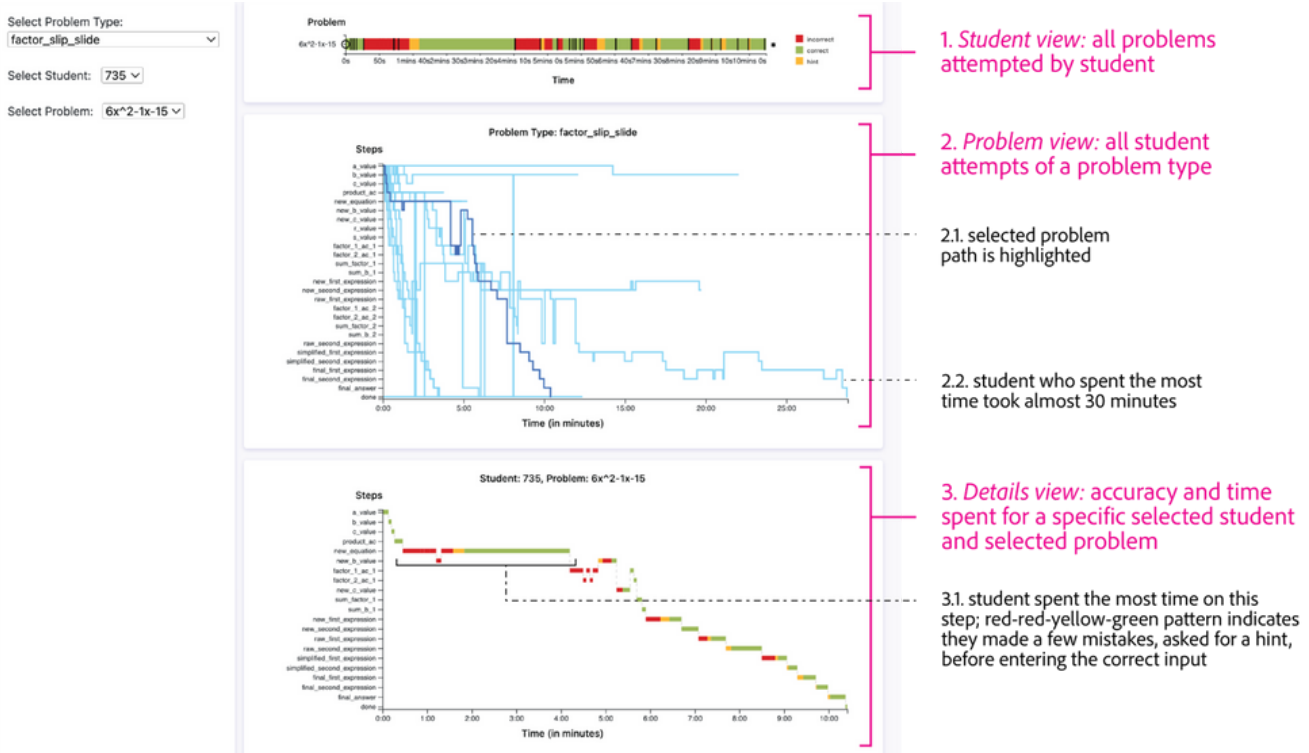
One of the suggestions that was raised during the NSF evaluation in Summer 2023 was to prioritize adult skill learning using VERA. To address this, in collaboration with the North Georgia Technical College (NGTC), we are conducting a second experimental study in Fall 2023. The goal of the study is to (1) investigate students' understanding of a "model" as a fundamental knowledge construct, and (2) "modeling" as a skill for students studying natural resource management. Note, to avoid confounding variables, at NGTC we will be deploying a non-personalized version of VERA. This contrasts with our study at Georgia Tech (mentioned earlier) which is aimed at assessing VERA's personalized learning system.

Breakdown of user activities in VERA (model construction, parameterization, and simulation) based on the identification of three general VERA users - constructor (Type 1), observer (Type 2), and full-explorer (Type 3). [Source: An S., (2023) Processes and Outcomes of Systems Thinking in an Interactive Modeling Environment. PhD Dissertation. Georgia Institute of Technology].

Visualization

We developed two data visualization systems as outputs of our research.

Vista (Figure below) is a web-based browser tool that helps instructors interactively analyze student data collected by Apprentice Tutors. Instructors deploying Apprentice Tutors in their classrooms can navigate to the “Analytics” tab of Apprentice Tutors to view visualizations of student performance, with detailed stepwise information about time spent, accuracy, and hints usage. The visualizations depict time and accuracy data for each step of a problem. Instructors can also select students or problem types to compare performance across multiple students or problems.



1. Student view: all problems attempted by student

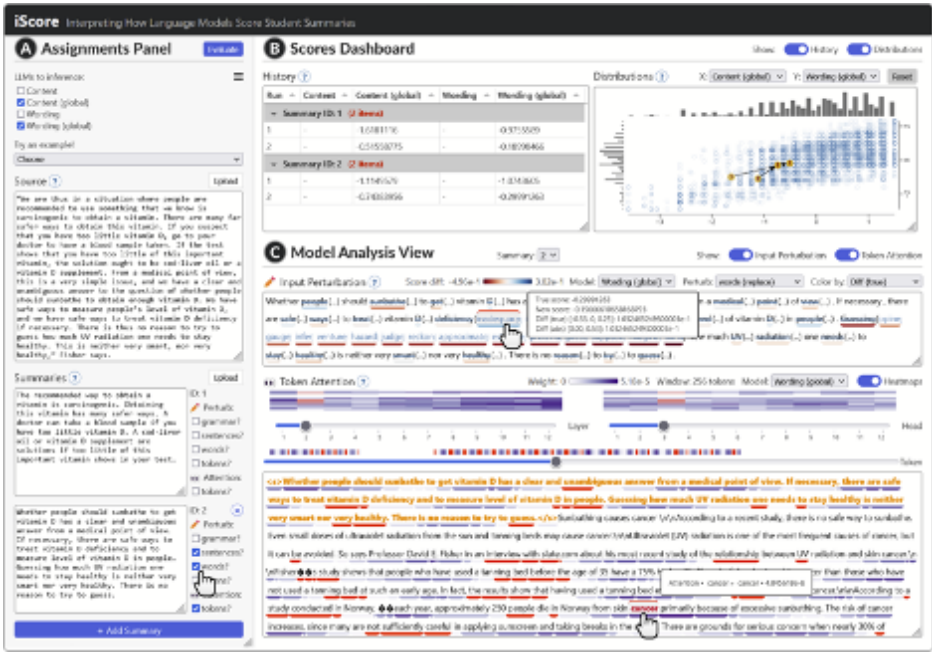
2. Problem view: all student attempts of a problem type

2.1. selected problem path is highlighted

2.2. student who spent the most time took almost 30 minutes

3. Details view: accuracy and time spent for a specific selected student and selected problem

3.1. student spent the most time on this step; red-red-yellow-green pattern indicates they made a few mistakes, asked for a hint, before entering the correct input



iScore (Figure on the left) is a web-based browser tool for researchers to upload, score, and compare multiple summaries of a source text simultaneously. iScore introduces a new workflow for comparing the language features that contribute to different LLM scores. First, users upload, score and can manually revise and re-score multiple source/summary pairs simultaneously. Then, users can visually track how scores change across revisions. Finally, users can analyze two LLM interpretability methods. Together, these views provide access to multiple summary comparison visualizations and several well-known LLM interpretability methods. Combining these visualizations and methods in a single visual interface broadly enables deeper analysis of LLM behavior that was previously time-consuming and difficult to perform.

AI-ALOE News & Events



AI-ALOE's During NSF AI Day and AI Hill Day in the Capitol: AI-ALOE's Director, Ashok Goel (Georgia Tech), along with Executive Committee Member Adie Shimandle (Technical College System of Georgia), showcased the Institute's progress in AI and adult education during the NSF AI Day and AI Hill Day at the Capitol, this month.



AI-ALOE turned 3 on November 1st: As our team members celebrates this milestone, we eagerly look forward to the exciting AI-based projects we'll bring in year three as we share our journey to transform online adult learning and online adult education!



AI-ALOE Members Participate in Inaugural Georgia Tech Showcase & Symposium: AI-ALOE members attended the Avant South Street Innovation Showcase in Tech Square's Coda courtyard and symposium Sept 28-29, 2023. This event functioned as a platform for various departments and Institutes to present their research, innovations, and demonstrations to a diverse audience that included students, staff, and the general public.



AI-ALOE Members Among Winners: Three Teams Secure Top Prizes in Mathematics Automated Scoring Challenge for The Nation's Report Card: From Vanderbilt University (Co-PI, Scott Crossley; Student Researchers, Langdon Holmes and Wesley Morris). More details: [AI-ALOE Members Among Winners: Three Teams Secure Top Prizes in Mathematics Automated Scoring Challenge for The Nation's Report Card.](#) For more, click [here](#).



AI-ALOE Welcomes New Members: We would like to formally introduce two additions to our management team: Moriah Ugi (picture on the left) will serve as our Assistant to the Executive, and Emma Green (picture on the right) is our newest Program Support Coordinator. Welcome!

AI-ALOE News & Events

Panel: How AI is empowering adult learning and online education

LLM: Training on MultiRC Data

Reference Models

- LLMs are provided with target answers provided by students
- Information from source is distilled as reference answer provided by ChatGPT.
- Models are provided candidate answers and references answers to predict whether the candidate answers are correct
- Models have no access to context beyond reference answers
- Correctness is determined by similarity between the target and the reference answer
- LLMs used
 - MPNet
 - BLEURT

LLM Input sequence:
[Candidate Answer] [sep token] [Reference Answer]

alaloe.org

on YouTube

AI-ALOE on Advancing Adult Learning with AI during ELAI GLOBAL 2023: AI-ALOE leaders and researchers Ashok Goel, Chris Dede, Scott Crossley, and Vrinda Nandan participated in a panel discussion titled: How AI Is Empowering Adult Learning and Online Education during the ELAI GLOBAL 2023 online conference hosted by Empowering Learners, Oct 23-25, 2023.



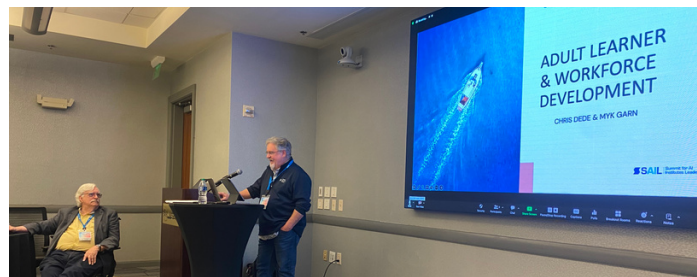
CNN Opinion: Ashok Goel on AI's Transformational Role in Education and Beyond: In a recent CNN opinion piece, Ashok Goel provides valuable insights into how artificial intelligence (AI) is shaping the new era of education. Drawing from his work on tools like Jill Watson, which not only has the potential to enhance human capabilities but also fosters creativity among students and professors, Goel highlights the transformative effects of AI. These include not only improved access to education but also the liberation of valuable time for essential tasks.

AI-ALOE AT SAIL 2023

MOMENTS AT THE 2023 SAIL SUMMIT:

AI-ALOE MEMBERS ENGAGE IN WORKSHOPS, PANELS, FOCUS GROUPS, AND EXPO DAY

AI-ALOE members played a key role in the Adult Learning Special Interest Group (AL-SIG) during SAIL. This group, led by our experts in adult learning from the AI Institute for Adult Learning and Online Education (AI-ALOE), focused on a crucial aspect for every AI Institute — understanding the needs of and implementing the best instructional practices for training stakeholders using AI tools and technologies. The session featured a special presentation by Nye Hodge, a Senior Research Analyst from the Atlanta Federal Reserve Bank's Center for Workforce and Economic Opportunity.



AI-ALOE was among the 25 NSF supported AI institutes on Thursday, October 26, for the first AI Institutes Expo Day at the Georgia Tech Hotel and Conference Center, in Atlanta, GA hosted by the AI Institutes Virtual Organization (AIVO).

The expo highlighted the institutes as each shared their cutting-edge AI research in various fields while discussing the current state of AI applications in the real world. This was followed by three-panel discussions with AI experts from the institutes on the Showcase Stage as well as a welcome from Georgia Tech President Ángel Cabrera



Collaboration & Outreach



OUTREACH

AI-ALOE's First Outreach Workshop: On November 2nd, AI-ALOE hosted its first inaugural outreach workshop. The workshop focused on the transformative potential of Artificial Intelligence (AI) in adult education and described how AI can enhance online education to help reskill, retrain, and upskill the next generation of American workers. This 3-hour event discussed cutting-edge developments in the fields of foundational and use-inspired AI with a special focus on recent work produced by researchers at AI-ALOE. Over 30 researchers participated in the workshop. AI-ALOE will host a second outreach workshop in spring of 2024 that will focus on adult learning.



AI-ALOE WEBINAR

SKILLS AS CURRENCY IN CONTEMPORARY LEARNING

INVITED SPEAKER

Shaun Moon


Instructure

VP of Product & Design Innovation

OCT 11 2023 | TIME: 12-1 PM

RSVP:
tinyurl.com/aialoeWebinar

AI-ALOE's Fall 2023 Webinars: AI-ALOE has welcomed two distinguished speakers for its webinars this Fall. Shaun Moon, VP of Production & Design Innovation at Instructure, gave a talk on skills as currency in contemporary learning during the October webinar. In November, Brandt Dargue, an Associate Technical Fellow, Senior Training Systems Integrator, and Learning Scientist at Boeing, discussed the current research on optimizing learning and continuously improving safety of flight. We are looking forward to having speakers from Google and in December and Wiley in January 2024.



AAAI Spring Symposium on Human-Like Learning: For the AAAI-24 Spring Symposium on Human-Like Learning, next year, AI-ALOE Researcher Chris MacLellan and Assistant Professor in the School of Interactive Computing at Georgia Tech is organizing a AAAI Spring Symposium on Human-Like Learning. His goal is to bring together researchers that have diverse perspectives on human-like learning and how we might create machine-learning systems with human-like capabilities (e.g., supporting online, incremental, continual, and data/energy efficient learning).

Team Member Spotlight



Meet AI-ALOE's Assistant Director for Outreach, **Scott Crossley**. He is a professor of special education at Vanderbilt University. His research focuses on using data science and artificial intelligence to explore cognition, with an emphasis on natural language processing.

SCOTT CROSSLEY



Building Academic Connections to Transform Education

By Scott Crossley

As the Director of Outreach for AI-ALOE, my primary role is to develop and build relationships with the academic and educational communities we serve. My work is all about ensuring that our mission at AI-ALOE is effectively communicated. To achieve this, I'm constantly organizing various events, such as workshops, symposiums, and listening sessions. These engagements help us engage with the people who share our interests and values, allowing us to grow and make a more significant impact in our field.

One of the flagship projects that my team and I are working on is the development of an intelligent textbook framework. This framework is designed based on the fundamental principles of reading skill development and comprehension. We believe that it has the potential to revolutionize the way we approach education and learning.



Speaking specifically about the project this project we're developing what we call "intelligent textbooks." These textbooks are not your typical static, ink-on-paper materials. Instead, they are adaptive, interactive, and intelligent. At Vanderbilt University, we've taken a significant step in this direction with the creation of Intelligent Texts for Enhanced Language Learning, or iTELL for short. This is a computational framework that has the capability to transform any machine-readable text into an interactive, intelligent format that can be accessed through a web application.

iTELL is rooted in established theories of reading comprehension, and it offers users a variety of features. Students using iTELL can write summaries about text and videos to showcase their understanding, and these summaries are automatically scored by large language models specifically trained for this purpose. The scores generated by these models provide qualitative feedback to students, helping them understand and improve their content and wording. iTELL also uses AI to generate concise questions for different sections of the text and videos, serving as a tool to assess comprehension. We are also working on integrating a structured think-aloud assistant into iTELL, powered by generative AI. This assistant can help users make inferences, provide relevant elaborations for the text, and connect the content to their personal experiences.



The feedback and insights derived from these AI integrations can be utilized in a myriad of ways, such as guiding learning, correcting misconceptions, reviewing missed topics, preparing for upcoming materials, drawing connections between the text and real-world experiences, and aiding in a deeper understanding of the material.

It's truly been an exciting journey so far, and I'm privileged to be part of groundbreaking work, like the iTELL project, in hopes of reaching the right audiences in meaningful and impactful ways on education and learning.

Student Highlight



**QIAOSI WANG
(CHELSEA)**

Meet **Qiaosi Wang**, also known as **Chelsea**, a Ph.D. candidate specializing in Human-Centered Computing at Georgia Tech, where she works under AI-ALOE Director Ashok Goel in the Design & Intelligence Lab (DILab) within the GVU Center. Chelsea's research focuses on the intersection of Human-AI interaction, Cognitive Science, and Computer Supported Cooperative Work (CSCW). Her work with AI-ALOE revolves around the Mutual Theory of Mind framework, which draws inspiration from the human innate ability to infer what's happening in others' minds (referred to as "Theory of Mind"). This framework aims to improve mutual understanding between humans and AI during their interactions.

The Mutual Theory of Mind Framework in Human-AI Interaction

By Qiaosi Wang

I'm a researcher at AI-ALOE, where I focus on improving human-AI interactions, cognitive science, and computer-supported cooperative work. My research aims to develop the Mutual Theory of Mind framework, which helps humans and AI better understand each other during their conversations.

Many researchers at AI-ALOE are dedicated to creating personalized and adaptive AI agents for online education. These AI agents play various roles, such as helping students understand the agent's abilities, assisting confused students, and reaching out to socially isolated students.

To do this effectively, AI agents need to understand the complexities of human thought, similar to how humans use Theory of Mind to make guesses about each other's thoughts based on words and actions. My work on the Mutual Theory of Mind framework guides the design of more personalized, ethical, and human-centered AI in online education.

My journey into this research began in 2019 when we deployed an AI agent named Jill Watson to answer questions about courses at Georgia Tech. We noticed that students had varying perceptions of Jill's capabilities, which led us to the idea of having Jill detect and clarify these perceptions. This concept led to our CHI 2021 paper, which showed that we could predict how students perceived Jill based on the length and sentiment of their questions. This work laid the foundation for the Mutual Theory of Mind, and we've been exploring human-AI interactions based on this framework ever since.

Outside of my research, I enjoy outdoor activities like hiking and bouldering to clear my mind, but sometimes my cat, Gouda, keeps me at home, which I don't mind at all!

MENTAL MODEL IN HUMAN-AI INTERACTION

SPEAKERS

- Qiaosi Wang (Chelsea), Ph.D. Candidate, Georgia Tech
- Sarah Wash, Ph.D. Candidate, Georgia Tech
- Zeina Khatami, Research Staff Member, IBM Research
- Mei Si, Associate Professor, NYU
- Ashok Goel, AI-ALOE Director, Professor, Georgia Tech
- Michael Muller, Research Staff Member, IBM Research
- Jeffrey Koppert, Distinguished Research Staff Member, IBM Research

Join us for a symposium exploring diverse mental models in human-AI interaction and more!

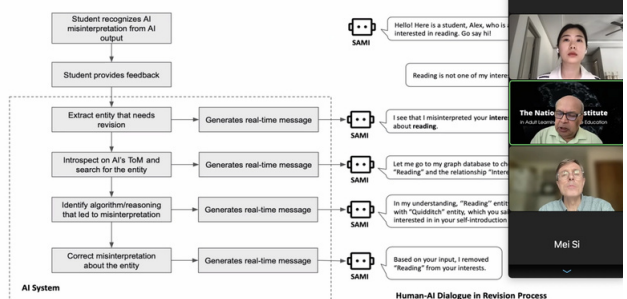
Date: Thursday, June 8, 2023
Time: 3pm to 6pm ET
Location: Virtual (Must Register)

[More information at aialoe.org](#)

AI-ALOE

The Mental Models in Human-AI Interaction symposium began with a presentation by Ashok Goel, the Director of AI-ALOE, who discussed the topic of how AI can self-explain and self-correct. The event continued with other topics relating to human-AI interactions. The symposium concluded with an open discussion about Large Language Models (LLMs), which are deep learning algorithms capable of performing various natural language processing (NLP) tasks. The discussion also touched on the concept of mental models and their potential applications in manipulation.

Agent Self-Models and Self-Corrections



Publications

AI-ALOE researchers have contributed more than 10 publications this quarter.

Bae, Y., Kim, J., & Kim, M. (2023). Clustering cognitive engagement changes in a longitudinal traced discussion data from an online course. In Damşa, C., Borge, M., Koh, E., & Worsley, M. (Eds.). *Proceedings of the 16th International Conference on Computer-Supported Collaborative Learning - CSCL 2023* (pp. 297-300). Montreal, Canada: International Society of the Learning Sciences.

Baffour, P., Saxberg, T., & Crossley, S. (2023). Analyzing Bias in Large Language Model Solutions for Assisted Writing Feedback Tools: Lessons from the Feedback Prize Competition Series. In *Proceedings of the 2023 BEA 18th Workshop on Innovative Use of NLP for Building Educational Applications*. (pp. 242–246). [10.18653/v1/2023.bea-1.21](https://doi.org/10.18653/v1/2023.bea-1.21)

Crossley, S. A., & Holmes, L. (2023). Assessing receptive vocabulary using state-of-the-art natural language processing techniques. *Journal of Second Language Studies*, 6(11), 1-28.

Crossley, S. A., Choi, J. S., Scherber, Y., & Lucka, M. (2023). Using large language models to develop readability formulas for educational settings. In N. Wang, G. Rebolledo-Mendez, V. Dimitrova, N. Matsuda, & O. C. Santos (Eds.), *Artificial Intelligence in Education (AIED)* (pp. 422–427). Springer Nature Switzerland.

Crossley, S. A., Cushing, S., Kyle, K., & Jarvis, S. (2023). Crowd-Sourcing Human Ratings of Linguistic Production. In M. Goldwater, F. Anggoro, B. Hayes & D. Ong (Eds.), In *Proceedings of the 45th Annual Conference of the Cognitive Science Society*. (pp. 1515-1520). Sydney, Australia: Cognitive Science Society.

Crossley, S. A., Wan, Q., Allen, L., & McNamara, D. S. (2023). Source inclusion in synthesis writing: An NLP approach to understanding argumentation, sourcing, and essay quality. *Reading and Writing*, 36, 1053–1083.

Holmes, L., Crossley, S. A., Morris, W., Sikka, H., & Trumbore, A. (2023). Deidentifying Student Writing with Rules and Transformers. In N. Wang, G. Rebolledo-Mendez, V. Dimitrova, N.

Holmes, L., Crossley, S. A., Sikka, H., & Morris, W. (in press). PIILO: An open-source system for personally identifiable information labeling and obfuscation. *Information and Learning Science*.

Hornback, A., Buckley, S., Kos, J., Bunin, S., An, S., Joyner, D., & Goel, A. (2023, July). A Scalable Architecture for Conducting A/B Experiments in Educational Settings. In *Proceedings of the Tenth ACM Conference on Learning@Scale* (pp. 373-377).

Kim, J., Haddadian, G., & Kim, M. (2023). An investigation of knowledge-based AI vs. human evaluation in academic summary evaluation: Similarities, dissimilarities, and being toward mutual understandings. In Blikstein, P., Van Aalst, J., Kizito, R., & Brennan, K. (Eds.). *Proceedings of the 17th International Conference of the Learning Sciences - ICLS 2023* (pp. 994-997). Montreal, Canada: International Society of the Learning Sciences.

Kim, M., Kim, N., Haddadian, G., & Heidari, A. (2023). A test of learning progress models using an AI-enabled knowledge representation system. In Blikstein, P., Van Aalst, J., Kizito, R., & Brennan, K. (Eds.). *Proceedings of the 17th International Conference of the Learning Sciences - ICLS 2023* (pp. 986-989). Montreal, Canada: International Society of the Learning Sciences.

Lawley, L., & MacLellan, C. (2023). VAL: Interactive task learning with GPT dialogue parsing. arXiv preprint arXiv:2310.01627.

**To see the full list of publications
on our website, click here**

Presentations

The AI-ALOE researchers have made significant contributions within their research areas, of over 13 presentations, invited talks, and panel discussions at the broader community beyond AI-ALOE.

Below you will find a selected list of our **External Presentations from AI-ALOE Researchers**:

Presenter	Presentation/Event	Date
Rahul Dass	Interactive Videos: Revisited and Beyond. Invited talk as part of AI-ALOE. Defense Chief Learning Officers Council (DCLOC)	Oct 2023
Ashok Goel	Teaching and Learning in the Age of AI EDULEARN 2023 Keynote	July 2023
Christopher MacLellan	Video games as a path to a contextualized cognitive science, or How to play 20 questions with nature and win CogSci-23 Workshop	July 2023
Scott Crossley	Better matching texts to readers using computationally informed readability formulas Victoria University of Wellington	July 2023
Scott Crossley	What is writing proficiency and how can we teach it better? Victoria University of Wellington	July 2023
Scott Crossley	Personalization in intelligent textbooks using large language models Pens & Pixels: Generative AI in Education, University of California Irvine	July 2023
Sanaz Ahmadzade Siyahrood & Michael Hoffmann	Revolutionizing Learning with Artificial Generative Futures Symposium	Sept 2023
Sanaz Ahmadzadeh Siyahrood & Michael Hoffmann	The best of technology, innovation, and creativity Avant South at Georgia Tech	Sept 2023
Michael Hoffmann	Harnessing the Force of AI Georgia Digital Government Summit	Oct 2023
Bae, Y., Kim, J., Haddadian, G., Davis, A., & Kim, M.	The impact of an AI-based educational tool, with a focus on technology acceptance and metacognitive awareness of adult learners The 2023 Association for Educational Communications and Technology (AECT) Conference, Orlando, FL	Oct 2023
Haddadian, G., Kim, J., Bae, Y., & Kim, M.	A Comprehensive Model of AI Literacy from a Developmental Perspective. The 2023 Association for Educational Communications and Technology (AECT) Conference, Orlando, FL.	Oct 2023
Kim, J., Bae, Y., Haddadian, G., Morris, W., Crossley, S., Holmes, L., Stravelakis, J., & Kim, M.	AI-augmented summarization: Impact on online adult learners' concept learning, discussion quality, and achievement The 2023 Association for Educational Communications and Technology (AECT) Conference, Orlando, FL.	Oct 2023
Kim, J., Bae, Y., Haddadian, G., & Kim, M.	Leveraging Machine Learning to Automatically Evaluate Cognitive Engagement in Asynchronous Online Discussions The 2023 Association for Educational Communications and Technology (AECT) Conference, Orlando, FL.	Oct 2023
John Kos	Systems Thinking and Inquiry-Based Learning with VERA Defense Chief Learning Officers Council (DCLOC)	Oct 2023

To see the full list of presentations on our website, [click here](#)

