



# THE AI-ALOE SPOTLIGHT



**Ashok Goel**  
Executive Director

As we head into spring, I'm excited to share the latest AI-ALOE Spotlight, jam-packed with updates and achievements.

First, a huge congratulations to AI-ALOE EC member Professor Ruth Kanfer for winning the prestigious Dunnette Prize for her work in industrial-organizational psychology. She'll be honored at the SIOP Annual Conference in April. We are very proud of Professor Kanfer and her accomplishments.

In this issue, we're highlighting the work of Alex Endert, an associate professor at Georgia Tech, who leads AI-ALOE research in Visual Analytics, improving AI's role in adult education. Also, we'd like to introduce two of our talented Ph.D. candidates, Adam Coscia and Grace Guo, who also contribute to the Visual Analytics Lab.

We also introduce new initiatives In this newsletter: the AI-ALOE Outreach Workshops, the AI-ALOE Research Fellows program, and the AI-ALOE Teaching Fellows Program. Our Outreach Workshops brought together experts to explore AI for adult learning. We're thrilled to welcome 27 Research Fellows and launch the Teaching Fellows Program for educators in Georgia.

On the research front, our teams are making great progress. The A4L team, for example, is analyzing data to show A4L's impact on learning. We're also improving technologies like SMART for better learning outcomes.

Exciting times ahead as we continue to push the boundaries of education with AI.

- Ashok Goel

## SPRING 2024 NEWSLETTER

### Table of Contents

#### Page 1

Director's Message

#### Page 2

A4L

#### Page 3

Apprentice Tutor Builder

#### Page 4

Interactive Videos (IVY)

#### Page 5

iTELL

#### Page 6

Jill Watson

#### 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

Collaboration & Outreach

#### Page 15

Team Member Spotlight

#### Page 16-17

Student Highlight

#### Page 18

Publications

#### Page 19

Presentations

### AI-ALOE Social Media



# A4L Technology Infrastructure

In Spring 2024, the A4L (AI Augmented Architecture for Adult Learning) team initiated a comprehensive data analysis, leveraging both recent and historical learning data from multiple AI tools and various data sources within courses at TCSG and Georgia Tech. The primary objective is to showcase the impact of A4L on learning. In doing so, we will leverage A4L's data engine to empower researchers with access to learner data, over multiple episodes, facilitating in-depth meso-learning analysis.

The initiative is focused on:

- Demonstrating A4L's impact on learning by granting researchers access to A4L's analytics data for meso-learning.
- Identifying new constraints and directions for enhancing A4L's architecture, design, and development.
- Defining future research goals and experiments for technologies like SAMI, VERA, JW, etc., specifically in the context of meso-learning.
- Developing visualizations and dashboards tailored to meet the needs of both teachers and learners.

A sub-team is dedicated to understanding SAMI's influence on learners' social presence and engagement, utilizing the Community of Inquiry framework as a guiding lens. Additionally, there is an exploration of equity considerations when integrating AI into the learning process, evaluating SAMI's value in terms of social, teaching, cognitive presence, and learner performance.

A JW focused sub-team is aiming to showcase how we can leverage JW data combined with other student data to explore impact of teaching presence on student cognitive engagement essentially on demographics (age, gender, race) that could lead to inclusive and equitable quality in education. We are also planning to develop a prototype of the analytics pipeline on A4L architecture platform.

A sub-team is also focused on understanding the opportunities for mesolearning in VERA. Specifically, we aim to demonstrate how student learning history, cognitive engagement, and demographics are important factors in how students interact with the VERA simulation platform and the various coaches that exist within VERA, with the ultimate goal of uncovering new mechanisms for enhancing and personalizing the learner experience. Much of this mesolearning work is made possible by A4L, and we aim to migrate the analytical components we develop to the infrastructure.

# Apprentice Tutor Builder

## Apprentice Tutors

- **Vital Prize:** The U.S. National Science Foundation (NSF), in partnership with the Bill & Melinda Gates Foundation, Schmidt Futures, and the Walton Family Foundation, held the Visionary Interdisciplinary Teams Advancing Learning (VITAL) Prize Challenge, with the goal of catalyzing a new generation of emerging technologies and innovations for K-12 learning. Apprentice Tutors & Apprentice Tutor Builder advanced into the semi-final round of this challenge and we received \$20,000 in research and development funding.
- **New Apprentice Tutors & Enhanced Tutor Functionality:** In the past few months we have created and deployed five new tutors within the TCSG ecosystem. These tutors include the following subjects: Logarithms, Quadratic Functions, Quadratic Equations, Exponential Equations, and Logarithmic Equations. We have also enhanced the Apprentice tutor capabilities to include improved hint strategies, a real time progress bar on the top of the tutor, and supplementary video / textbook resources mapped to every tutor. These enhancements stem directly from the user focus groups conducted mid-2023, and we anticipate these changes will have a positive effect on overall tutor likability leading to great tutor adoption / retention.
- **Adult-learning Experimentation:** We are planning to run some randomized controlled trials (RCTs) centered around adult learning outcomes to see how the Apprentice Tutors can best support students. We plan to implement faded, adaptive scaffolding within each problem type. As the student approaches mastery, we will show fewer and fewer sub-steps, encouraging students to fully compile the problem-solving process. We hypothesize this will improve retention and build more independent, confident learners.
- **Integration of HTN planners with Apprentice Tutors:** We are shifting the core infrastructure of the tutors from rete algorithm to hierarchical task networks (HTNs). We will leverage HTNs for robust model tracing, in addition to enabling precise knowledge tracing at multiple hierarchy levels. By accurately estimating students' mastery over different methods, HTNs will significantly enhance personalized learning experiences.

## Apprentice Tutor Builder

- **VAL for Hints:** We are exploring ways to connect our Verbal Apprentice Learner into the Apprentice Tutor Builder. One method built last semester is generating declarative hints based on the HTN acquired by ATB. By using Chat-GPT to verbalize the methods and subtasks in the HTN, we can build robust, yet safe and accurate hints and instructions.
- **Language-Based Tutor Generation:** We are exploring how to use LLMs such as ChatGPT to let ATB users create tutors by providing a natural language description of the tutor they want. A preliminary prototype shows that this approach can produce interface elements relevant to the given description with much of the scaffolding in place. Users can then further edit and customize the tutor to their needs. This feature helps achieve our goal of ATB being usable by those with no specialized technical or design training.

# Interactive Videos (IVY)

We have continued to iterate over our proposed infrastructure for a novel Interactive Video or “Ivy” coach by extending the work done by the former “In-Video” tutor project [1-3] using Georgia Tech’s Knowledge-based AI (KBAI) course from the Online Master’s in Computer Science (OMSCS) program. First, we identified and based our theoretical foundations for learning and change using the ICAP (Interactive, Constructive, Active and Passive) theory of cognitive engagement [4]. This has enabled us to better understand how different modes of cognitive engagement, demonstrated by an online learner, relate to various changes in cognitive-knowledge processes. In turn, this has helped us formulate several research issues and hypothesis regarding how we envision an Ivy coach would help a learner foster more active to constructive and eventually interactive forms of learning.

However, prior to helping any learner, an Ivy coach must first have some level of understanding of the skill that a learner is being taught. Using the TMK (Task-Method-Knowledge) language [5], we developed two preliminary TMK models to conceptualize how an Ivy coach would solve problems pertaining to three skills - Semantic Networks, Generate and Test, and Means-End Analysis from the KBAI course. We are currently developing additional TMK models to further investigate their viability as a generalized method for modeling skills and problem-solving.

Although we ran into several challenges when implementing the Ivy coach such as the unavailability of the In-Video tutor codebase and Udacity deprecating OMSCS courses from its platform, we developed an initial Ivy prototype by creating a “Udacity-like” clone using next.js. This prototype gave us increased flexibility in terms of (1) integrating Ivy within EdStem – where the OMSCS program is currently hosted, (2) adding as many quizzes and videos as needed, and (3) enabling us with direct access to collect learner and coach-related data from EdStem for future user-experimental studies and analysis.

Finally, by leveraging TMK models as a conceptual representation for understanding skills, we are exploring how to generate explanations via reasoning using various Generative AI methods for different types of learner-Ivy coach interactions. We recently presented our progress during the Spring 2024 AI-ALOE Foundational AI and Use-Inspired AI weekly meeting series and are encouraged by the positive feedback and suggestions provided by the AI-ALOE community.

Constructed response items, in which students are prompted to provide an open-ended response to a question and are commonly used to assess reading comprehension they encourage active processing of information and to improve learning, relative to simply reading. However, scoring of constructed responses can take considerable time and resources. Automatic generation of questions and automatic scoring of constructed responses to those questions could have a significant impact on the ability to deploy these types of constructed response questions at scale.

We have integrated automatically generated constructed responses into the intelligent Texts for Enhanced Lifelong Learning (iTELL) framework to support interactive learning for a variety of complex subjects to provide students with opportunities to better comprehend material. iTELL uses GPT-3.5 to generate questions and correct constructed responses for each language chunk. During use, the student submits their answer to the constructed response and the answer is scored by two LLMs fine-tuned for the purpose.

We used an iTELL deployment of an introduction to computer programming textbook to collect usage data from 98 participants who produced 2,733 constructed responses. Participants were invited to provide feedback on the constructed response task while they were responding to the questions by clicking on a thumbs up or thumbs down button and selecting up to three tags to explain why they provided that feedback. We collected data from 38 participants of which 75.4% responses were positive (see Figure 1).

Participants also completed an outtake survey, and their responses were used to determine the efficacy of the automatic short constructed response scoring feature in iTELL. The majority of of respondents (81.7%) endorsed agreement or strong agreement to the statement that the short answer tasks helped them to improve their learning. Similarly, 84.1% of respondents agreed or strongly agreed that the constructed response task was easy to work with and 86.6% of respondents endorsed agreement or strong agreement to the statement that the questions were relevant to the subsection. A majority of 73.1% of participants also agreed or strongly agreed that the model accurately scored their responses (see Figure 2).

Participants broadly approved of the constructed responses in iTELL, with more than four fifths of users expressing agreement or strong agreement with the statement that the task helped them improve their learning. The majority of the negative feedback had to do with the accuracy of the feedback, which provides next steps for the development of the constructed response integrated into iTELL.

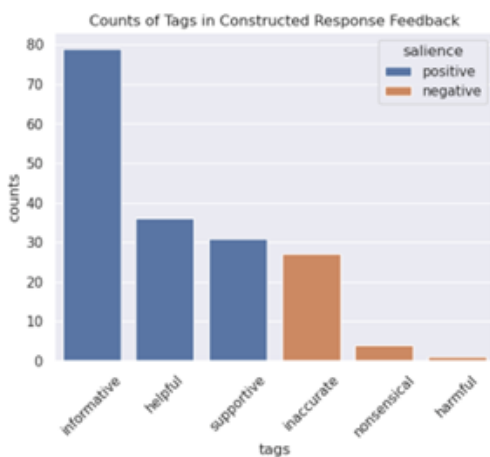


Figure 1: Counts of Feedback Tags

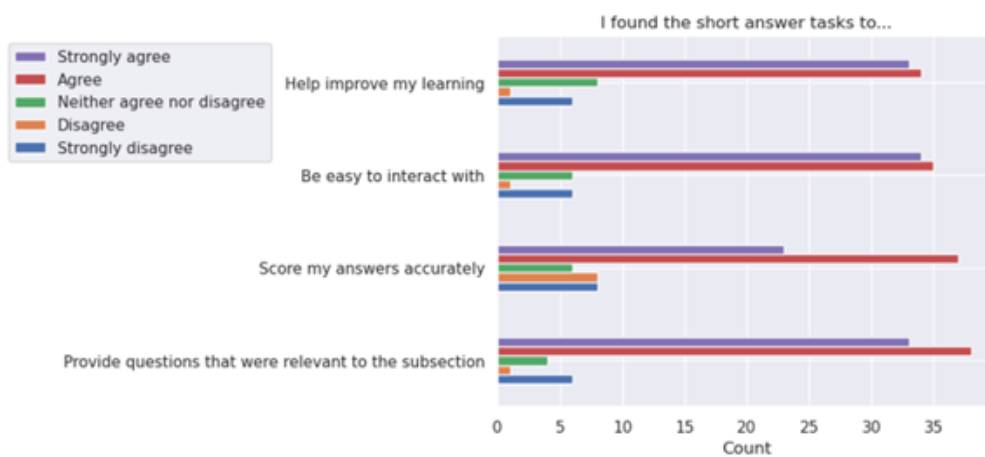


Figure 1: Feedback on Constructed Response Task

# Jill Watson

Jill Watson is an AI teaching assistant that engages students in extended conversations about educational materials such as textbooks, video transcripts, presentation slides, and class syllabi.

In Spring 2024, the Jill Watson team completed three major changes: generating synthetic QA datasets, improving context retrieval, and automated evaluations.

1. Synthetic CQA dataset generation: Creating manual datasets for testing Jill's performance before each deployment cycle is both cost and time intensive. We have now developed a synthetic QA generation module, which uses Google's Vertex AI to synthetically generate relevant question-answer pairs for each chunk in each document.
2. Improved Information Retrieval: Jill uses the context retrieval module to retrieve the top n most relevant texts in an agent's knowledge base and then submits those texts to ChatGPT for response generation grounded in truth. If the context retrieved is incorrect, Jill could provide incorrect but reasonable responses to students. Based on experiments conducted in Fall 2023, Jill now uses an updated DPR model and has been equipped with a cross-encoder to re-rank the output from the DPR model. This has improved MRR and Recall@5 by up to 15% on various manually generated datasets. These results were also confirmed by using synthetically generated QA datasets.
3. Automated evaluation: Jill uses a textual entailment test to confirm the consistency of generated responses, which could double as an automated evaluation module for pre-deployment testing. We have tested advanced prompting strategies.

In Fall 2023, we conducted an A/B experiment in OMSCS KBAI class at Georgia Tech to determine the impact of Jill in improving teaching, cognitive and social presence in classrooms. Our analysis of student responses to a survey based on the Community of Inquiry guidelines shows a small, positive, statistically significant effect on the Design & Organization dimension of Teaching Presence as well as Social Presence. We also notice some statistically significant improvement in exam results, but no overall consistent difference in performance in the two groups on the other assignments in the course.

In Spring 2024, we have deployed Jill in 5 courses – two OMSCS courses at Georgia Tech, English Composition and Rhetoric course at Wiregrass College, Leadership course at Columbus Tech and Intro to Programming on EdX.

# Participatory Design

The goal of the team “Participatory Design for Human Well-being” is to ensure that all AI technologies developed by the AI-ALOE Institute will be designed in a way that takes the well-being of users, unintended users, and stakeholders into account. To achieve our project goal, the team has performed the following major activities over the past four months.

- 1. Well-Being Impact Analysis (WIA) of the Apprentice Tutor Technology:** The PD team prepared a Well-being Impact Analysis (WIA) report of the Apprentice Tutor technology, developed by Professor Christopher J. MacLellan and his team. This report covers the internal ethical analysis, stakeholder engagement process initiated through the focus group meeting, and presents a list of actionable items based on feedback received from participants who attended the meeting on April 28, 2023. The stakeholder engagement section discusses both positive and negative impacts of the Apprentice Tutor on users' well-being and the barriers that limit its usage. Included is also a list of action items that could or should be implemented to address the negative impacts on well-being that we observed.
- 2. Identified Challenges and Strategies for Overcoming Them:**
  - **Equity Issues in Focus Group Meetings - NSF requested to specify effects of AI technologies on various demographic groups** - The main question is whether AI technologies affect various demographic groups equally or differently. For the PD team this poses the challenge to align the feedback provided by participants in focus groups to the specific demographics of those individuals who provide feedback about well-being effects. The question is how corresponding data can be collected in a non-intrusive manner and how the privacy of focus group participants can be protected. The PD team asked experts and practitioners at the AI-ALOE Institute for their support to address this challenge: Dr. Aileen Reid, Alex Endert, and Jacob Dallas-Main. Together with these experts, the PD started a conversation on these challenges that is still ongoing. A proposal for a strategy is expected for the second quarter.
  - **A new name for our team to replace the notion of "Participatory Design"** - Each of the concepts “participatory design” and “co-design” has been used to signify design processes that actively engage and involve end-users throughout the design process. “Throughout” means that potential users are involved already in the generation of design ideas in communication that attempts to identify their needs, interests, and values or the problem that should be addressed with a technology or product. This is not what the PD team is doing. Our team realizes responsible AI design by engaging users and other stakeholders only after a first version of an AI technology has been deployed and used. For this reason, the team has suggested adopting the term “user-engaged design” going forward to describe its work. We suggest the following name: **User-engaged Design for Human Well-being**. An investigation into the usage of “user-engaged design” in the academic literature revealed its mention in three Web of Science publications, two ACM Digital Library entries, and seventy Google Scholar references. However, only two Google Scholar resources provided an explanation of how the term has been used. All or most of the others are using the term only to illuminate other terms such as co-design. Since the term “user-engaged design” is not yet defined in the literature, our team proposes the following definition: User-engaged design is a design process that engages and involves users of a technology and other stakeholders after a first version or a first prototype of this technology has been developed. The goal of user-engaged design is to garner feedback from potential users and stakeholders to shape the further development of this technology or to stop this development.

# 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 Quarter 1 2024, we made significant progress in fortifying SAMI's metacognitive abilities in error localization, correction, and explanation. SAMI's explanation abilities were deployed in live classes. Using a Task-Method-Knowledge model, SAMI can explain its own architecture and capabilities. We hope that these improvements will help users to build trust in the agent through building a more accurate mental model of how SAMI functions, or theory of SAMI's mind. We also expanded how SAMI encourages students to interact with one another: prompting them to "click to say hi" to their matches. This provides users with a low-stakes method of acknowledging a match and reaching out to them. Not only does this provide us with another metric for measuring student interaction with SAMI, but it also allows students to easily observe their peers' interactions. We hope that through increased interaction with this feature, students will increase their interaction with matches overall.

We have continued to finetune the prompts SAMI uses with ChatGPT. This includes prompts for named entity recognition (NER) and for framing SAMI's responses containing match suggestions for students. As a result, SAMI's responses have more accurate NER results and remain natural, engaging, and diverse. We have continued explorations into SAMI's effects on social presence and the potential for matching based on new dimensions, including personality and features beyond those we currently identify.

In January 2024, we deployed SAMI in two classes in the OMSCS program and one in-person CS course. Both online courses (with about 500 and 1200 students) were deployed on Ed discussion.



## SMART Research Papers Submitted to ICLS 2024

The SMART research aimed to pursue dual goals: one was to hone AI-driven assessment techniques, and the other was to conduct design research on the impact of AI-augmented concept learning based on the multi-cohort data gathered from English and Biology courses at the Technical College System of Georgia (TCSG), offered since Spring 2022. Consequently, we submitted five research papers to the International Conference for the Learning Sciences (ICLS) 2024. Two papers were related to the former goal. More specifically, a paper examined how AI tracked learner changes in their concept learning by comparing knowledge-based and large-language model (LLM)-based AI approaches. This study was conducted in collaboration with Dr. Scott Crossley's team at Vanderbilt University. The other paper tested generative-AI-driven concept learning assessment through the comparison with human-generated evaluation in preparation for the next-generation evaluation techniques for SMART. For the latter goals, two studies investigated learner mental model changes in SMART-based assignments (i.e., micro-level learning) and the relationships between SMART learning experience and the subsequent learning activities (i.e., meso-level learning) in English and Biology courses, respectively. Additionally, working with Dr. Michael Hoffman's team at Georgia Tech, we submitted an intensive case study of a participatory design research we had conducted during the past two years with the aim of enhancing SMART technology with a focus on well-being in education. Moreover, we gathered data from Fall 2023, during which we conducted A/B experiments: a quasi-experiment in English classes and a Randomized Control Trial (RCT) experiment in a biology class. These experiments aimed to compare learner engagement and performance in two different SMART versions. We are now analyzing the data.

## Release of SMART Block 30 in January 2024 and W.R.I.T.E. Prototype Design

We launched the latest version of SMART, named Block 30, in January 2024 to enhance personalization through multimodal formative feedback. The new version utilizes generative AI, ChatGPT, to create a concept library that enables learners to access concise definitions of key concepts they may have missed in their previous attempts. It also provides examples illustrating the relationships between these concepts in propositional statements. Additionally, improved SMART data visualization now highlights individual learners' behavior patterns, including their revision history, resulting levels of understanding, and time spent reviewing specific feedback components. Benchmark information is also provided, displaying the class's average engagement and performance levels. The Block 30 was introduced to TCSG faculty members during the showcase meeting on January 19th, 2024, and it has been deployed in the TCSG classes in Spring 2024.

We have also initiated a new project to develop the Writing and Reasoning Intelligent Tutor for Education (W.R.I.T.E.) system. This system builds on the advantages of the current SMART but fully utilizes generative AI models to offer a highly interactive learning experience between AI and human agents, along with more advanced personalized feedback information. To achieve this, we have designed an application framework for W.R.I.T.E. and prototyped the user experience and interface.

# VERA

VERA (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 Spring 2024, we refined 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. We launched both our AI coaches on the live VERA site, and we conducted a study to assess the personalized version of VERA in three live Georgia Tech classrooms: an introductory ecology course in Fall 2023 (N > 80 students), an advanced ecology course in Spring 2024 (N > 35 students), and a cognitive science course in Spring 2024 (N > 50 students). We are currently conducting data analysis on these datasets, and we are using both quantitative and qualitative methods to demonstrate how effective a personalized feedback system can be designed for ill-defined, inquiry-based learning environments.

In Summer 2023, we deployed a version of VERA into a Georgia Tech cognitive science classroom with the ability to request feedback from a coach, and our purpose in doing so was to determine when learners requested feedback the most. We recently finished our analysis of this data, and we found that around 70 percent of coach requests were made while learners were simulating their conceptual models. This highlights the fact that simulations are learners' opportunities to evaluate their models, and it is where they may need the most support if their evaluation is not as expected. This will guide our future work in developing personalized coaches.

We have also collaborated with the North Georgia Technical College (NGTC) to conduct an experimental study on adult learning. This study aimed to achieve two goals:

1. Investigating 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 deployed 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. We are conducting data analysis for this study, too.

Finally, we are continuing development on our Ask-TMK project, which uses a TMK (Task-Method-Knowledge) model to support self-explanation in VERA. Currently, we are developing a prototype that aims to interpret and generate explanations with reasoning regarding questions pertaining to learners' conceptual models and real-time simulation results.

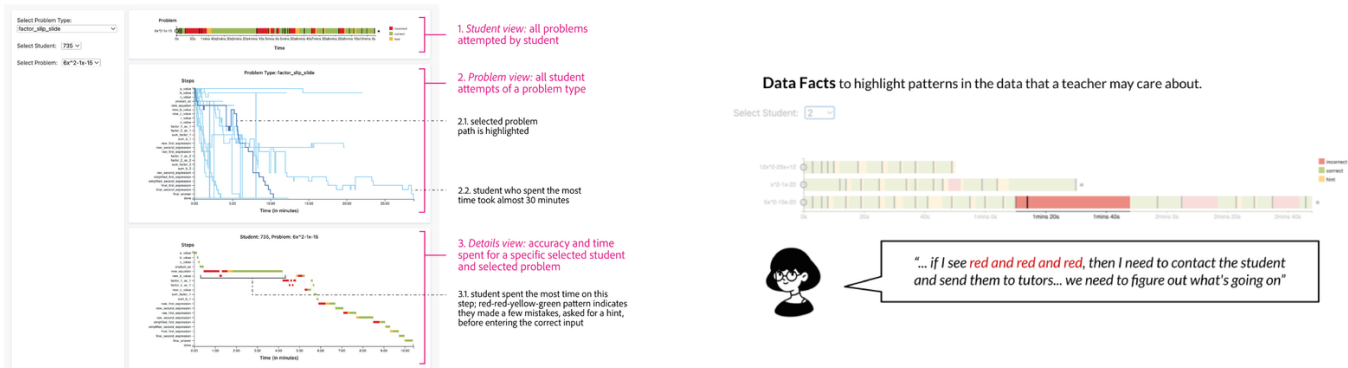
# Visualization

We developed two data visualization systems to aid in personalization during deployment, and increasing trust in AI prior to deployment.

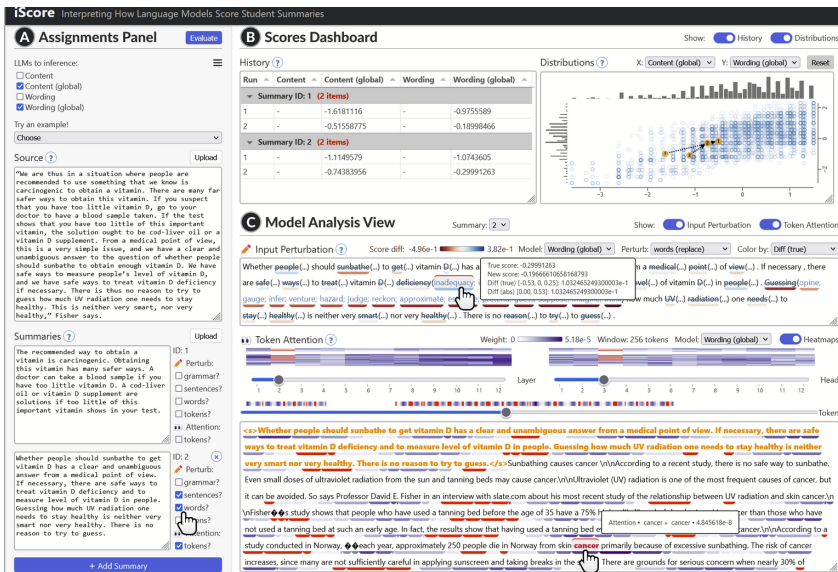
VisTA (Figure below) is a web-based browser tool that helps instructors interactively analyze student data collected by Apprentice Tutors from student usage. In a focus group study with five teachers who have deployed VisTA in their classroom, we found that teachers often wanted to highlight particular patterns in student problem solving behavior. For example, in the quote below (bottom right), one teacher said that if she saw “red and red and red” in a visualization (signifying groups of students continuously getting a question incorrect), it would indicate that she needs to follow up with the student in question. Our current study implements an interactive pattern query function for teachers to surface such insights automatically. Moving forward, we plan to conduct a deployment study comparing just the visualizations (bottom left) to a system where both visualizations and data facts are available to understand how such “auto-insights” can aid in personalization.

VisTA

VisTA + data facts



Score (Figure below) is a web-based browser tool for AI-ALOE researchers to upload, score, and visually compare multiple summaries of a source text simultaneously. The summaries are written by students in AI-ALOE's iTELL and then evaluated by large language models (LLMs). Comparing language features (semantics, syntax, etc.) that contribute to LLM scores enables a deeper analysis of LLM behavior that was previously time-consuming and difficult to perform.



iScore also addresses major ethical concerns around deploying LLMs in educational tools, ensuring fairness and transparency. iScore was accepted to the ACM Intelligence User Interfaces (UI) conference (top in the field) and will be presented in March 2024. We are currently working on a follow-up visualization dashboard for teachers that use iTELL in their classes, helping them track summary scores and make decisions, adapting and expanding on our best designs from iScore.

# AI-ALOE News & Events (1)



## AI-ALOE Showcase: Save the Date

Join us for the inaugural National AI Institute for Adult Learning and Online Education (AI-ALOE) showcase! Explore cutting-edge advancements in AI for learning and education at this interactive event featuring a series of engaging activities:

- **3-Minute Madness:** Lightning talks highlighting innovative AI-ALOE projects. (11:00 am - 12:00 pm)
- **Posters and Demonstrations:** Experience firsthand the latest developments through interactive displays, accompanied by a light lunch. (12:00 pm - 1:00 pm)
- **Panel Discussion:** Delve deeper into the intersection of AI and education with insights from Georgia Tech leaders and researchers. (1:00 pm - 2:00 pm)

Don't miss this opportunity to connect with leaders and innovators in the field of AI-enhanced learning. Save the date and be part of the conversation!

## AI-ALOE Monthly Webinars

AI-ALOE's monthly webinar features the following distinguished researchers, scientists, and experts in AI and education during this past quarter of the year:

11/8/2023: **Brandt Dargue**, Learning Scientist at Boeing, gave a talk on research in advancing aerospace training at Boeing to optimize learning and continuously improve safety of flight.

12/13/ 2023: **Jeff Rubenstein**, Education and AI Solutions Manager at Google, discussed the potential value(s), applications, and challenges of AI in education, and how Google is looking to address these challenges; and some ways to help think through the coming evolutions in AI.

1/24/2024: **Andrew Jones**, Director of AI and MLOps at Wiley, presented the frameworks and use cases of generative AI at Wiley.

2/14/2024: **Norm Vaughan**, Professor in the Department of Education at Mount Royal University, Canada, provided an overview to the Community of Inquiry framework and discussed how AI can be used to scale up the framework.



# AI-ALOE News & Events (2)

## AI-ALOE External Advisory Board Meeting



AI-ALOE convened its virtual [External Advisory Board \(EAB\)](#) on December 8, 2023. Ashok Goel began the meeting by presenting an overview of the EAB members and sharing progress that the institute has made since the last EAB meeting. Progress included iTELL being deployed along with personalization in model learning. The logic models and issue hypothesis trees were shared as well and it was noted that these may evolve and change as the work evolves.

The focus topics for the meeting centered around generative AI and LLMs, with brief projects updates from each team followed by more detailed presentations by the Architecture for Learning and Responsible AI team. The EAB members were proactive and engaging throughout the meeting, asking questions about the institute's progress and providing recommendations on the institute's future research directions.

## AI-ALOE Researchers Host PII Data Detection Competition



AI-ALOE's co-PI, Professor Scott Crossley and his team at Vanderbilt University partnered with the Learning Agency Lab to host a competition focused on developing automated techniques for detecting and removing personally identifiable information (PII) from educational data, specifically in student writing. The primary goal is to create models that can effectively identify and eliminate PII, ultimately reducing the cost of releasing educational datasets. Successful automation in this area would benefit learning science research and facilitate the development of educational tools by enabling access to large public educational datasets.

The competition aims to support researchers and industry professionals in leveraging the potential of such datasets for the creation of effective tools and interventions to support teachers and students. There is \$60,000 in prize money for [the PII data detection competition](#) and the entry deadline is April 16, 2024.

# AI-ALOE News & Events (3)



## AI-ALOE Faculty Awarded Prestigious SIOP Dunnette Prize

Ruth Kanfer, Professor in the School of Psychology at Georgia Tech, and an Executive Committee member of AI-ALOE, has won the Dunnette Prize from the Society for Industrial and Organizational Psychology (SIOP), for her transformative theorizing and research on how individual differences in abilities, skills and motivation combine to influence workplace learning and job performance. She will officially receive the honor and deliver an invited address at the SIOP Annual Conference.

[Learn more](#)



## Barbara Oakley Received 2023 McGraw Prize

AI-ALOE recognizes Barbara Oakley, one of AI-ALOE External Advisory Board (EAB) members, for receiving the 2023 McGraw Prize in recognition of her contributions to lifelong learning. Oakley's dedication and expertise truly aligns with the goals and work of AI-ALOE.

In a recent statement, Oakley emphasized the potential impact of integrating “learning how to learn” into AI-ALOE’s personalized learning approach. She highlighted that this addition could generate measurable and positive results, particularly benefiting the most disadvantaged learners whom our initiative aims to support.

Oakley has also made learning easier for millions of people from all walks of life by revealing how insights from neuroscience can improve our understanding of challenging concepts in math and engineering—and countless other fields. To hear more from Barbara Oakley on this subject, you can [watch this video](#).



## Steven Ferguson Featured in CXOTech Magazine

A new article featured in [CXOTech Magazine](#) by Steven Ferguson, CIO of the Technical College System of Georgia and AI-ALOE Executive Committee Member, explores the significance of driving digital innovation in organizations. Ferguson shares his journey and emphasizes the need for embracing emerging technologies, fostering a culture of innovation, and forming strategic partnerships. He highlights recent trends like AI, IoT, cloud computing, and cybersecurity while advocating for a balanced approach to adopting new technologies. Ferguson stresses the transformative impact of digital innovation on industries and urges leaders to adapt to stay competitive.

# Collaboration & Outreach


## AI-ALOE Outreach Workshops and Research Fellows



To aid in AL-ALOE's role as a nexus for the fields of adult learning, online education, and artificial intelligence, the Institute hosted two invitational research outreach workshops in December 2023 and January 2024. The first workshop centered on cutting-edge developments in both foundational and use-inspired AI, particularly for enhancing online occupationally-focused learning for adults. The second centered on specific ways AI can improve online learning to help the entry, upskilling, and reskilling of the next generation of American workers. The goal of the workshops was to foster innovation and collaboration by briefing small groups (15-20) of experts in specific topics about the work of AI-ALOE, then offering participants the opportunity to become affiliated AI-ALOE Research Fellows.

The workshops resulted in AI-ALOE adding a very distinguished group of [27 Research Fellows in Artificial Intelligence and in Online Learning](#). They have some opportunities to participate in our events, seminars, and workshops. Additionally, a dedicated Slack channel serves the nexus for the research fellows, fostering an environment conducive to idea exchange, project exploration, and collaborative research endeavors in AI and education. Discussions with that group about mutual topics of interest are ongoing.

## AI-ALOE Teaching Fellows Program



AI-ALOE is proud to offer [the Teaching Fellows Program](#), an exceptional opportunity for educators within the Technical College System of Georgia (TCSG) passionate about leveraging AI technologies to transform teaching and learning experiences. This fellowship empowers instructors to participate in cutting-edge research studies at AI-ALOE, collaborating with our esteemed researchers and scientists to integrate AI solutions into their curricula.

We will call for applications among instructors at TCSG soon. We are excited to collaborate with selected teaching fellows to enhance learning with AI technologies, cultivating a vibrant community of educators engaged in transforming teaching and learning with innovative use of AI.

# Team Member Spotlight



**Alex Endert**

Alex Endert is an associate professor in the School of Interactive Computing at the Georgia Institute of Technology. He directs the [Visual Analytics Lab](#), which focuses on designing and studying how interactive visual tools assist individuals in comprehending data and AI. The lab frequently tests these advancements across various domains, including intelligence analysis, cybersecurity, learning, decision-making, manufacturing safety, among others. Generous support for their work comes from sponsors including NSF, DOD, DHS, DARPA, DOE, and industry partners.

Endert received his Ph.D. in Computer Science from Virginia Tech in 2012. In 2013, his work on Semantic Interaction was awarded the IEEE VGTC VPG Pioneers Group Doctoral Dissertation Award, and the Virginia Tech Computer Science Best Dissertation Award. In 2018, he received a CAREER award from the National Science Foundation for his work on visual analytics by demonstration.

## Visual Analytics to Understand, Personalize, and Trust AI in Adult Education

Our research in AI-ALOE aims to enable individuals involved in education to better interpret, understand, and monitor the use of AI for various tasks. The goal is to support three primary user groups involved in AI-ALOE: teachers, learners, and AI researchers who develop the AI agents for various purposes. Each user group faces specific challenges as their tasks and goals differ.

For example, through focus groups and iterative design practices, we have discovered that teachers are primarily interested in understanding how learners utilize the AI agents in the context of the learning goals and pedagogical structure they envision. We have developed a visualization (called VisTA) to allow teachers to track how students utilize intelligent tutors to learn specific skills. Our findings indicate that beyond simple tracking and monitoring performance, typically resulting in assigning grades, teachers are more interested in understanding learning behavior. For example, visually exploring the data in aggregate reveals gaps in knowledge across a group of learners, leading to adjustments in materials or lessons to ensure students better understand the material.



Our research has also enabled AI developers to test, improve, and ultimately calibrate their trust and confidence in specific AI agents prior to deployment. Specifically, we have developed visual analytic tools to illuminate errors in how LLMs auto-grade text summaries written by students. These tools (iScore and KnowledgeVis) have revealed biases and other linguistic errors, resulting in improvements in the design of LLMs before deployment. Capabilities like this aid in ensuring the ethical deployment and use of AIs for educational purposes in AI-ALOE.



# Student Highlight




**Adam Coscia**

Adam Coscia is a Human-Centered Computing Ph.D. student at Georgia Tech and works closely with Alex Endert as a member of the [Visual Analytics Lab](#). His work focuses on designing and developing visual analytics tools to aid individuals in comprehending complex data, mitigating biases, and facilitating decision-making processes. Currently, his primary focus is assisting NLP researchers and engineers in analyzing and interpreting insights gained from large language models (LLMs) throughout their training, inference, and comparative performance stages. Beyond academia, Coscia has contributed to real-world applications, partnering with scientists engaged in field and mission operations. Together, they've crafted visual analytics solutions to support deep ocean expeditions and enhance scientist-guided spacecraft autonomy. Coscia's journey has involved collaboration with professionals from various institutions, including Vanderbilt, Georgia State, Tufts, Emory, Caltech, ArtCenter, Monterey Bay Aquarium Research Institute (MBARI), and NASA Jet Propulsion Laboratory (JPL).

## What is your primary research area(s) with AI-ALOE?

My research goal with AI-ALOE is to enable end-to-end evaluation of large language models (LLMs) deployed in novel educational technologies. The learning community has begun integrating LLMs like ChatGPT into adaptive learning tools for improving adult education. The benefits of LLMs are numerous: enabling new methods to generate educational content, powering new tools for automatically evaluating student work, providing new interfaces for assisting students in the learning process, and more. However, multiple stakeholders, from linguistics researchers to learning engineers and even teachers, are concerned with safely deploying LLM-powered technology in critical learning environments. Engaging in a human-centered design process, I have begun designing, developing, and deploying actionable and interpretable data visualizations that help stakeholders understand and validate LLM performance through an interactive and scalable interface with LLM data, ultimately making the decisions that our AI produces more transparent and responsible.



## What motivates and guides your research at AI-ALOE?

My research addresses ethical concerns with deploying LLMs in educational tools, ensuring fairness and transparency. By developing visualization tools, I aim to enable stakeholders to understand and interpret LLMs' performance effectively.

## How did you become interested in your research field?

My interest in explaining complex data led me to pursue a PhD in Human-Centered Computing at Georgia Tech. I design visual analytics tools to help people understand and explain AI systems deployed in critical environments like classrooms.

## Do you have a favorite hobby outside of research?

Lots! I love sports. I play baseball, softball, golf, and tennis. I enjoy bouldering, hiking, swimming, and snowboarding in the winter. I go for lots of runs, and I recently started bodybuilding! I play lots of video games. I want to try my hand at game design and development someday. I also enjoy reading fantasy novels like "The Lord of the Rings." I cook every day, trying new recipes and techniques all the time! From time to time, I enjoy trying different art projects such as making pottery, painting, and drawing. And of course, I love doing all of these activities with my friends and family.

## Could you share an interesting and enjoyable tidbit about yourself?

Fun facts— I taught snowboarding professionally for 5 years in Eastern Pennsylvania, and I've interned at NASA's Jet Propulsion Laboratory (JPL) twice!

## RECENT PUBLICATION

Adam Coscia, Langdon Holmes, Wesley Morris, Joon Suh Choi, Scott Crossley, and Alex Endert. 2024. iScore: Visual Analytics for Interpreting How Language Models Automatically Score Summaries. *In 29th International Conference on Intelligent User Interfaces (IUI '24), March 18–21, 2024, Greenville, SC, USA*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3640543.3645142>

# Student Highlight

in



Grace Guo

Grace Guo is a Human-Centered Computing Ph.D. student at Georgia Tech and works closely with Alex Endert as a member of the [Visual Analytics Lab](#). She is interested in work at the intersection of AI, visualization and human-computer interaction. Gup's current research focuses on developing visual analytics tools for explainable AI and causal inference, particularly in the domains of adult education and healthcare analytics. I am honored to be a recipient of the [2023 IBM PhD Fellowship](#).

## What is your primary research area(s) with AI-ALOE?

My primary research areas within AI-ALOE revolve around the intersection of educational data analytics and visualization development. Specifically, I examine how students interact with intelligent tutors being developed here at AI-ALOE. My work explores how visualizations of student problem-solving processes can help instructors understand student performance and enhance the overall learning experience, emphasizing personalized and adaptive approaches. This research also extends to the broader applications of AI in educational settings, aiming to harness the potential of these technologies for transformative impacts on teaching and learning methodologies.

## What motivates and guides your research at AI-ALOE?

I am motivated by an interest and belief in the progress of educational technologies. The recent advancements in AI can be leveraged to provide a greater diversity of learning material tailored to individual students. At the same time, teachers remain essential in providing guidance, social support, and experienced pedagogy. My research seeks to provide teachers with insight into how their students are using intelligent tutors, with the goal of facilitating a seamless integration of the strengths of both AI and human instruction.

## How did you become interested in your research field?

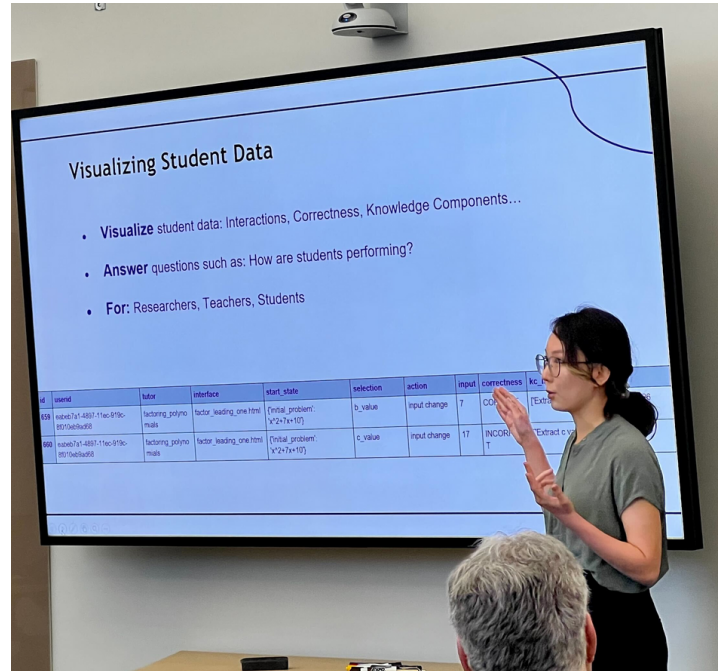
Visualization has long been used to enhance human cognition in many fields, helping us identify and communicate key ideas in data that would otherwise be obscured. Combined with the advancing landscape of educational technologies, I believe that visual analytics can intersect with pedagogy in novel and impactful ways.

## Do you have a favorite hobby outside of research?

Cooking and baking, with different levels of skill or results.

## Could you share an interesting and enjoyable tidbit about yourself?

I grew up in the tropics, where there is (to this day) a great diversity of local fauna. My house has been broken into twice: once by a macaque monkey, and another time by a civet.



## RECENT PUBLICATION

Guo, G., Karavani, E., Endert, A., & Kwon, B. C. (2023). Causalvis: Visualizations for causal inference. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3544548.3581236>

# Publications

- Baffour, P., & Crossley, S. A. (in press). Advances in automating feedback for argumentative writing: Feedback Prize as a case study. In J. Wilson & M. Shermis (Eds). *Routledge International Handbook of Automated Essay Evaluation*. New York: Routledge.
- Bondie, R., Mancenido, Z., Adams, H., Dede, C. (2024). Exploring data visualization in mixed reality simulations to measure teacher responsiveness. In: Bourguet, M.L., Krüger, J.M., Pedrosa, D., Dengel, A., Peña-Rios, A., Richter, J. (eds) *Immersive Learning Research Network. iLRN 2023*, pp. 173-181. *Communications in Computer and Information Science*, vol 1904. Springer, Cham.
- Coscia, A., & Endert, A. (2023). KnowledgeVIS: Interpreting language models by comparing fill-in-the-blank prompts. *IEEE Transactions on Visualization and Computer Graphics (IEEE TVCG)*.
- Crossley, S. A., Heintz, A., Choi, J., Batchelor, J., Karimi, M., & Malatinszky, A. (2023). A large-scaled corpus for assessing text readability. *Behavior Research Methods*, 55, 491–507.
- Crossley, S. A., Tywoniw, R., & Choi, J. S. (in press). The tool for automatic measurement of morphological information (TAMMI). *Behavior Research Methods*.
- Dede, C., & Lidwell, W. (2023). Developing a next-generation model for massive digital learning. *Education Sciences*, 13(8), 845.
- Goel, A., Dede, C., Garn, M., & Ou, C. (in press) AI-ALOE: AI for reskilling, upskilling and workforce Development. *AI Magazine*.
- Goel, A., Nandan, V., Gregori, E., An, S., & Rugaber, S. (2024). Explanation as question answering based on user guides. In S. Tulli & D. W. Aha (Eds.), *Explainable Agency in Artificial Intelligence: Research and Practice*. Routledge.
- Goel, A., Ou, C. (in press) Introduction to the special issue on national AI Institutes. *AI Magazine*.
- Kim, M. K., Gaul, C. J., Bundrage, C. N., & Madathany, R. J. (2023). Technology supported reading comprehension: a design research of the student mental model analyzer for research and teaching (SMART) technology. *Interactive Learning Environments*, 31(3), 1377-1401.
- Kos, J., Eaton, K., Zhang, S., Dass, R., Buckley, S., An, S., & Goel, A. (2023). Using analytics on student created data to content validate pedagogical tools. *arXiv preprint arXiv:2312.06871*.
- Lawley, L., & MacLellan, C. J. (2024). VAL: Interactive task learning with GPT dialog parsing. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing System*.
- Lee, J., Nandan, V., Sikka, H., Rugaber, S., & Goel, A. (2023). Designing a communication bridge between communities: Participatory design for a question-answering AI agent. *arXiv preprint arXiv:2308.00813*.
- Lyndgaard, S. F., & Kanfer, R. (2024). Interpersonal, intrapersonal, and cognitive tactics: A thematic analysis of adults' 21st century learning management. *Journal of the Learning Sciences*, 1-37.
- Morris, W., Crossley, S., Holmes, L., Ou, Chaohua, Dascalu, M., & McNamara, D. (in press). Formative feedback on student-authored summaries in intelligent textbooks using large language models. *International Journal of Artificial Intelligence in Education*.
- Ou, C., & Joyner, D. (2023). Seven years of online project-based learning at scale. *International Journal on Innovations in Online Education*, 7(1).
- Taneja, K., & Goel, A. (2024). Machine teaching for building modular AI agents based on zero-shot learners. *arXiv preprint arXiv:2401.05467*.
- Zhang, Q. (2023). Understanding human-AI teaming dynamics through gaming environments. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 19(1), 440-443.

[Visit AI-ALOE on Google Scholar for a complete list of publications](#)

# Presentations

Presenter(s)	Presentation/Event	Date
Ashok Goel, Scott Crossley, Chris Dede, & Vrinda Nandan	<a href="#"><u>How AI Is Empowering Adult Learning and Online Education / Empower Learners for the Age of AI (ELAI) 2023</u></a>	October 2023
Ashok Goel	The Frontiers of AI / <a href="#"><u>NSF Podcast</u></a>	November 2023
Scott Crossley & Landon Holmes	Word Surprisal Predicts Language Learner Proficiency / Annual Meeting of the Society for Computation in Psychology (SCiP), San Francisco, CA.	November 2023
Scott Crossley, J.S.Choi, & K.S. McCarthy	Improving Text Simplification in a Real-World Context: Leveraging the Automatic Readability Tool for English (ARTE)	November 2023
Sanaz Ahmadzadeh Siahrood	Participatory Design For Human Well-Being / <a href="#"><u>The Center for Integrative Research in Computing and Learning Sciences (CIRCLS) 2023 - Shaping AI and Emerging Technologies to Empower Learning Communities</u></a>	November 2023
Ashok Goel	Jill Watson / Defense Chief Learning Officers Council	January 2024
Adam Coscia	<a href="#"><u>iScore: Visual Analytics for Interpreting How Language Models Automatically Score Summaries</u></a> / Georgia Tech CoC Graduate Poster Symposium	January 2024
Scott Crossley	From tokens to transformers: Applications in text complexity / Learning Informatics Lab: Winter Colloquium. University of Minnesota	January 2024
Adam Coscia	iScore: Visual Analytics for Interpreting How Language Models Automatically Score Summaries / Georgia Tech Career, Research, and Innovation Development Conference (CRIDC) 2024	February 2024