



Letter from the Director

Ashok Goel

Executive Director, NSF AI-ALOE Institute



We have entered Year Five of AI-ALOE, and the progress made across our partner institutions and research teams has been significant. Instructor participation in the Technical College System of Georgia (TCSG) continues to rise, and more than 4,500 students used AI-ALOE learning tools in 310 classes last year. Our ecosystem of technologies, including *Apprentice Tutors*, *iTELL*, *Ivy*, *Jill Watson*, *SAMI*, *SMART*, *VERA*, and *A4L*, is now supporting adult learners at the TCSG, Georgia Tech, and Georgia State University in ways that are meaningful and measurable.

This year, our work in healthcare education remains especially important. High failure and withdrawal rates in key courses continue to limit the number of students who can move forward in nursing and allied health programs at a time when workforce shortages are growing. Early pilots at Southeastern Technical College, including drug-calculation support through the *Apprentice Tutors*, indicate that AI-supported learning can play a productive role. We expect deeper opportunities for evaluation when the tools are fully integrated in Fall 2026.

I should also note our growing work on education in AI (in addition to AI in education). We use a variety of tools - *Ivy*, *Jill Watson*, *SAMI*, and *VERA* - to foster education in AI in Georgia Tech's Online Master of Science in Computer Science program (OMSCS). The OMSCS is an archetypal example of using AI to teach AI for reskilling, upskilling, and workforce development. Last year, over 10,000 OMSCS students used at least one AI-ALOE tool. This year, we will also introduce *A4L* in OMSCS classes so that the AI tools can talk with one another.

Our community has also grown. I am delighted to welcome Dr. Carolyn Fonseca as AI-ALOE's new Program Manager. This newsletter highlights several colleagues whose contributions strengthen our work. Charmain Alston supports the Institute through dedicated grants administration. Rachel Sikes, who has extensive experience in critical care nursing and nursing education, serves as an AI-ALOE Teaching Fellow and applies our tools in pharmacology instruction. Rebecca Storey is one of Graduate Research Assistants and contributes important research on the motivational and social aspects of AI adoption among adult learners.

As we move through Year Five, AI-ALOE will continue to advance its technologies, strengthen the data infrastructure, and deepen partnerships across education and workforce domains. Our mission remains clear: develop AI tools that help adult learners succeed and support the preparation of professionals in fields where they are most needed. We look forward to the work ahead.

Page 1 | [Letter from the Director](#)

Page 2 | [Table of Contents](#)

Page 3 | [A4L](#)

Page 4 | [A4L Pipeline](#)

Page 5 | [Apprentice Tutors](#)

Page 6 | [Hallucination Suppression](#)

Page 7 | [Interactive Videos \(IVY\)](#)

Page 8 | [iTELL](#)

Page 9 | [Jill Watson](#)

Page 10 | [Machine Teaching](#)

Page 11 | [Mutual Theory of Mind \(MToM\)](#)

Page 12 | [SAMI](#)

Page 13 | [Self-Explanation](#)

Page 14 | [SMART](#)

Page 15 | [Tech Well-Being Assessment](#)

Page 16 | [VERA](#)

Page 17 | [Visualization](#)

Page 18 | [Staff Spotlight](#)

Page 19 | [Member Spotlights](#)

Page 20 | [Conferences and Talks](#)

Page 21 | [Research and Innovation](#)

Page 22 | [Featured Events](#)

Page 23 | [Upcoming Talks](#)

Page 24 | [Publications](#)

Page 25 | [Publications Cont.](#)

Page 26 | [Acknowledgment](#)

Table of Contents

🔍 Navigate quickly by clicking on a page number/title.



Follow AI-ALOE

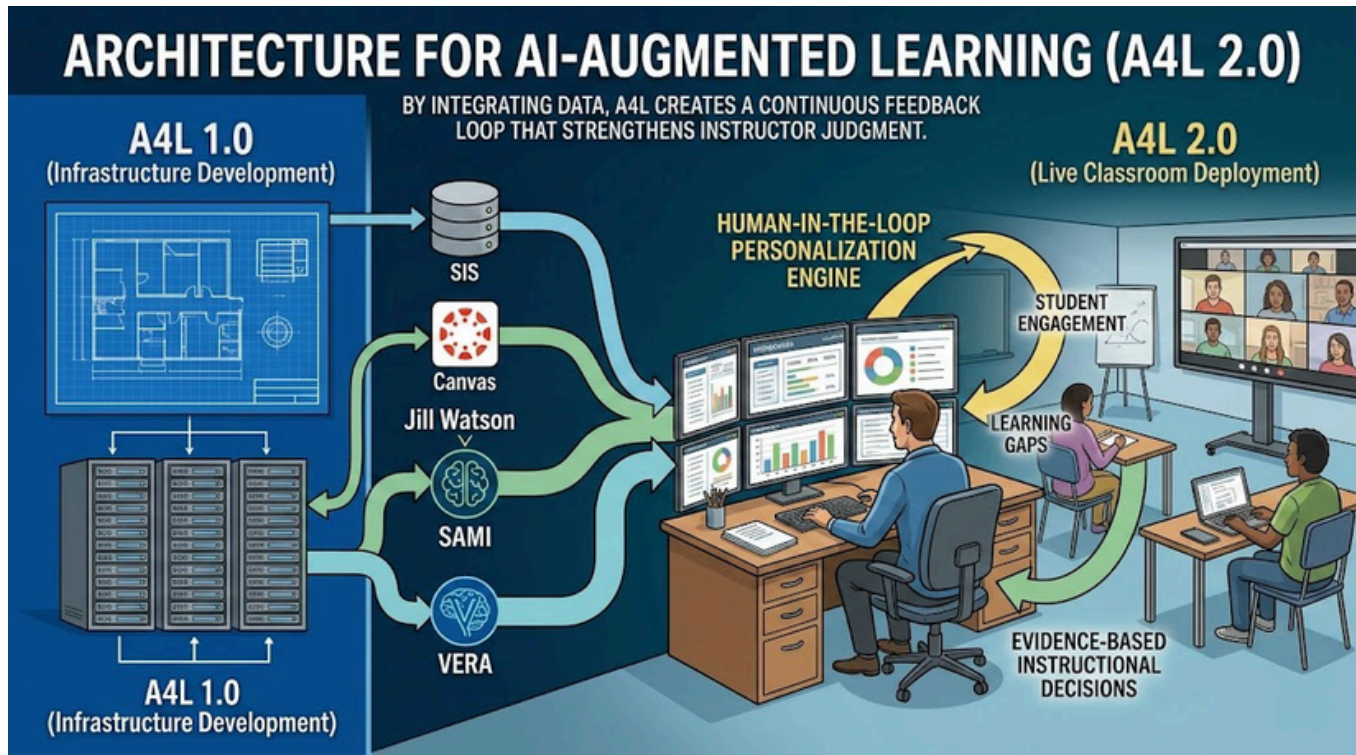


A4L Technology Infrastructure

A4L 2.0 Moves into the Classroom

Transforming Fragmented Data into Instructional Insight

The **Architecture for AI-Augmented Learning (A4L)** has moved from infrastructure to live deployment. A4L 2.0 aggregates data from SIS, Canvas, and AI tools (Jill Watson, SAMI, VERA) to give instructors a continuous, human-in-the-loop view of adult online learners.



From Activity Logs to Learner Trajectories

In large asynchronous courses, A4L 2.0 creates longitudinal learner representations that help instructors identify at-risk students, distinguish engagement from true mastery, and tailor interventions based on persistence patterns rather than clicks.

From Monitoring to Reasoning

Faculty co-design sessions show a shift from simple monitoring ("Who logged in?") to deeper reasoning ("What does this learner need next?").

Impact

By integrating multiple AI tools into one pipeline, A4L strengthens instructor decision-making and supports course completion and workforce readiness at scale.

A4L's Pipeline

A4L Pipeline Developed in Cooperation with 1EdTech

Overview

A4L 2.0 is an interoperable data pipeline designed to integrate SIS, LMS, and AI tool data (Jill Watson, SAMI, VERA, IVY) using 1EdTech open standards—Edu-API, LTI, and Caliper Analytics. It supports longitudinal research on how AI tools impact adult learning and strengthens AI-ALOE's sustainability as we enter year five.

Key Features

- Standards-based (LTI, Edu-API, Caliper)
- Real-time interoperable data ecosystem
- Unified learner and course context
- Shared analytics layer enabling cross-tool personalization and insights



Tim Couper, 1EdTech Chief Architect & Associate CTO presents

Why It Matters

Higher education institutions struggle to organize and use data effectively. A4L 2.0 provides a unified, standards-based infrastructure that reduces silos, improves decision-making, and supports scalable AI research across courses and tools.



1EdTech offered technical training to members of AI-ALOE

Deployment Progress

Over 12 months, the pipeline was deployed in Georgia Tech's AWS environment. The team finalized data governance, mapped SIS and LMS data to Edu-API, mapped AI tool events to Caliper, and validated secure data exchange via sFTP and the Asynchronous Bridge.

Current Experiments

- OMSCS 7637 KBAI: 500 students using JW, SAMI, VERA
- CS 4635/7637 KBAI: 150 students using JW, SAMI, VERA

Experiments will validate full-semester data ingestion, event capture, and analytics.

What's Next

Hands-on training for GA Tech developers (March 5–6, 2026)

- Launch of Interoperable Data Pipelines Course
- Introduction of A4L 2.0 to other AIVO institutes
- Development of A4L 2.0 reference architecture as a global 1EdTech standard

Apprentice Tutors

Apprentice Tutors Boosts Personalized Support for Adult Learners

The **Apprentice Tutors** platform now offers adaptive, fully online tutoring across the Technical College System of Georgia (TCSG), serving adult learners with an average age of 25.7 and growing demand in nursing and workforce programs.

High-Demand Program Focus

Originally used for college algebra, the platform now primarily supports nursing, which accounts for nearly half of all tutor interactions and most new users.

Enhanced Personalization and Scaffolding

Recent updates include:

- Flexible tutoring behavior via an upgraded HTN model
- Improved learner readiness estimates through Bayesian knowledge tracing
- Adaptive scaffolding that targets unmastered skills and reduces cognitive load



LLM Tutoring Limits

Understanding the Limits of LLM-Tutoring.
TutorGym evaluation shows LLMs struggle to spot student errors and provide hints.

Understanding the Limits of LLM-Tutoring.
TutorGym evaluation shows LLMs struggle to spot student errors and provide hints.

Supporting Growth

Supporting Student and Researcher Growth.
Empowering PhDs and postdocs for dissertations, new directions, and professional advancement.

Supporting Student and Researcher Growth.
Empowering PhDs and postdocs for dissertations, new directions, and professional advancement.

Instructor Partnerships

Partnerships With Instructors.
Ongoing collaboration with TCSG nursing instructors shapes tutor design, laying foundation for machine teaching.

Partnerships With Instructors.
Ongoing collaboration with TCSG nursing instructors shapes tutor design, laying foundation for machine teaching.

TutorGym:

Accelerating Innovation

TutorGym provides a unified environment for testing AI tutors and simulated learners, enabling rapid evaluation of instructional designs and research on student modeling.

Note: These insights show that adult learners engage more when support is available only when they choose to use it.

Key Findings From Deployments

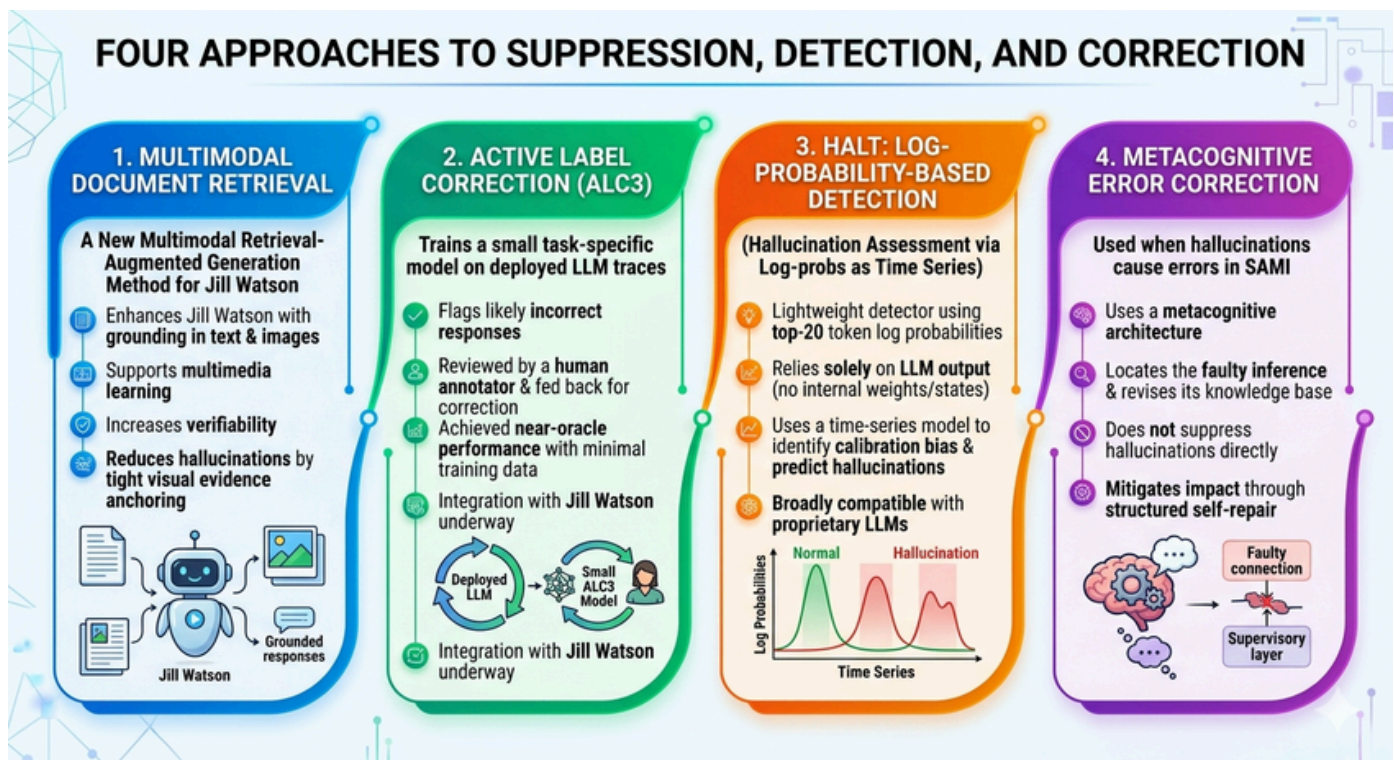
Category	Result
Timeframe	Spring, Summer, and Fall 2025
Course Sections Supported	256
Adult Learners Reached	1,000+
Study Type	A/B study on scaffolding approaches
Key Finding 1	On-demand scaffolding increased adoption by 50% compared with full scaffolding
Key Finding 2	Learners using on-demand scaffolding solved more problems and worked more efficiently

Hallucination Suppression

Advancing Hallucination Suppression and Correction in AI Systems

Why Hallucination Reduction Matters

Teachers and instructional designers frequently hesitate to adopt LLM-powered tools because of the risk of incorrect or unverifiable information. For adult learners, who often depend on accurate, timely support—the consequences of hallucination can be particularly disruptive. Improving the fidelity and transparency of AI responses is essential for building trust and usability in learning contexts.



Impact and Next Steps

These techniques illustrate promising pathways for reducing and correcting hallucinations across AI-ALOE technologies. By combining multimodal grounding, targeted correction, efficient detection, and metacognitive self-repair, the project is laying the foundation for more reliable AI partners in adult learning.

Interactive Videos (IVY)

IVY: Neuro-Symbolic Coaching for Procedural Skills

Ivy is an AI coaching system that helps adult learners master procedural skills by combining symbolic Task–Method–Knowledge (TMK) models with large language models (LLMs). Unlike unconstrained LLMs, Ivy explains what to do, how each step works, and why it matters. In Spring 2026, over 120 learners in Georgia Tech's Knowledge-Based AI course used Ivy to turn passive video watching into structured, explanation-driven learning.

Neuro-Symbolic Architecture

Ivy separates symbolic control from LLM generation. TMK models define goals, causal transitions, and hierarchical structure, while the LLM generates natural-language explanations constrained by these models. A four-stage pipeline ensures explanations remain causal, compositional, and reliable. Validated across 23 skills and 125 instructor questions, Ivy outperforms baseline methods in structural fidelity.

Scaling With Text-to-Model (TTM) Workflows

An AI-assisted text-to-model workflow enables rapid TMK model construction across entire courses, supporting scalable procedural modeling (under review at Learning @ Scale 2026).

Early Classroom Insights

In its first deployment, 102 learners generated 400+ procedural questions. Bloom's Taxonomy analysis shows Ivy handles more higher-order "how" and "why" questions than Jill Watson or forwarded inquiries, highlighting its potential for misconception detection and adaptive scaffolding.

Strategic Expansion

Ivy's growth is supported by grants (Spencer Foundation, NSF RITEL) and partnerships in workforce training. Engagement across AI-ALOE, including programming, AI literacy, and nursing education, is guiding Ivy's next generation to support diverse procedural skills and learning contexts.

Watch: Ivy – Interactive AI-Powered Videos for Deeper Learning



Click to watch short overview video. You'll learn:

1. How Ivy works and supports active learning
2. Its TMK + AI reasoning pipeline
3. Classroom and course deployments that enhance adult learning

Featuring AI-ALOE Members: Rahul Dass, Arpit Khandelwal, and Jeanette Luu discussing Ivy's vision, knowledge base, and impact on learning.

iTELL Expands Deployment Across Education and Industry:

Moving from Research-Centered development into broader use across higher education institutions and workplace training environments

The **Intelligent Texts for Enhanced Lifelong Learning (iTELL)** uses large language models trained to score learner work and generate feedback that helps users identify misconceptions, strengthen understanding and connect material to real-world contexts. The platform is grounded in established research on reading comprehension and adult learning.

Reading Support Through STAIRS

A core feature of the platform is STAIRS, or Strategic Thinking and Interactive Reading Support. STAIRS prompts readers to explain and elaborate on key ideas, helping them use evidence-based reading strategies associated with deeper understanding. The tool encourages users to monitor their thinking and develop skills that support lifelong learning. In addition to STAIRS, the system offers summary writing, cloze testing and constructed response items. Each task is automatically evaluated, giving learners immediate feedback they can use to improve comprehension and writing.

Personalized Learning for Adults

iTELL includes adaptive interventions designed for adult learners who often balance school, work and family obligations. The platform supports self-paced learning and offers optional, detailed feedback when users want additional help. Personalization draws on reading behavior, semantic analysis and Bayesian scoring. The system adjusts requirements for summary writing based on course difficulty, population performance and individual history. When learners struggle, STAIRS triggers a targeted re-reading activity focused on specific sections of the text.

Growing Use Across Courses and Workforce Programs

The platform is being used in online Python programming courses at Georgia Tech, research methods courses at Middle Georgia University and other STEM and vocational programs. Learners represent a range of ages, backgrounds and professional experiences.

Through the national I-Corps program, the iTELL team is also working with industry partners, including Chevron and Wolters Kluwer, to explore how the system can support workplace training and upskilling.

Partner Feedback Drives Improvements

Input from instructors and industry partners has guided improvements to the platform. Partners have assisted with text selection, validation of AI-generated questions and refinements to dashboards and feedback features. Their input has led to new quiz modules, updates to the preprocessing pipeline, interactive programming sandboxes and revisions to constructed response feedback.

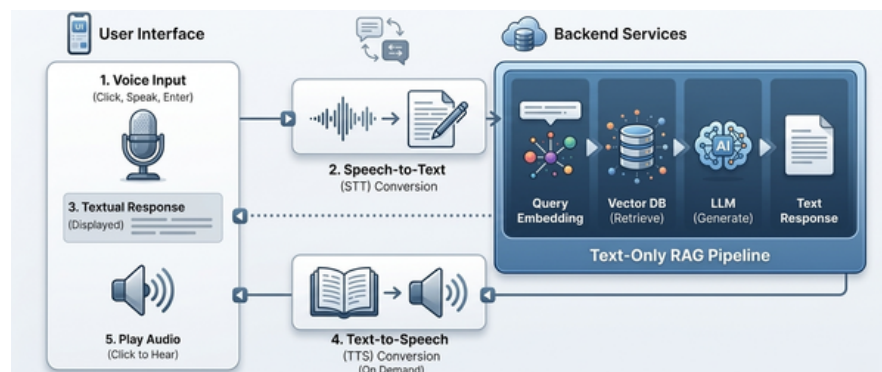
Research Output and Team Development

Recent project work has resulted in five conference publications and a book chapter. The platform has supported Ph.D. students, master's students, undergraduate researchers and technical staff. A former research associate recently accepted a role as a senior data analyst at The Learning Agency.

Jill Watson

Jill Watson: Building the Next Generation AI Learning Partner

The Jill Watson project is transitioning from a classroom chatbot into a scalable, multimodal cognitive partner that supports learners across online, hybrid, and workforce training environments. Built on instructor-approved course materials, Jill Watson helps adult learners access reliable, contextual support anytime—without increasing instructor workload. Recent work centers on multimodal reasoning, accessibility, and infrastructure that enables large-scale deployment.



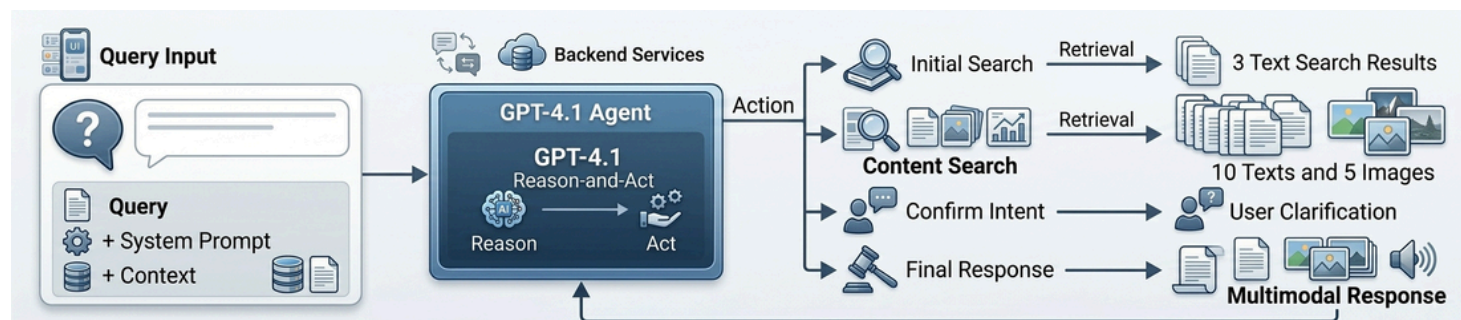
Multimodal Conversations With MuDoC

A major milestone is MuDoC, a new multimodal version of Jill Watson that blends text and images into grounded conversational explanations. MuDoC retrieves relevant diagrams, figures, and visuals from course documents, helping learners understand complex STEM ideas with clearer, more contextualized answers.

In a randomized study, MuDoC outperformed both a text-only chatbot and a textbook search tool on post-test learning gains. Learners rated MuDoC as more engaging, easier to navigate, and especially helpful for visualizing difficult concepts.

Voice Interaction for Real-World Learning

To support time-fragmented, on-the-go study, Jill Watson now offers push-to-talk voice interaction. Learners can ask questions verbally and receive spoken responses, a major benefit for those who struggle with reading, typing, or who study during commutes. Voice support is available in 100+ languages, expanding accessibility across diverse learning contexts.



Instructor-Deployable and Scalable

Jill Watson is now instructor-deployable, enabling faculty to upload course materials through a simple web interface and generate a course-specific assistant in under an hour, no engineering required. It runs on elastic cloud infrastructure and integrates with Canvas, Brightspace, Blackboard, and Moodle via LTI.

Interactive eTextbooks and Personalized Support

Embedded in digital textbooks, Jill Watson lets learners ask questions, request summaries, and receive feedback while reading. Personalization adapts responses to preferred levels of detail and reasoning style without retraining models.

Advancing Jill Watson as a Cognitive Partner

Upcoming features include Socratic mode for guided inquiry, assessment mode for conversational evaluation, full-duplex voice interaction, and a mobile app. Researchers are also studying how AI influences teaching presence, engagement, and learning assessment.



Watch: Jill Watson, The AI Teaching Assistant Transforming Education

Machine Teaching

Apprentice Tutors Boosts Personalized Support for Adult Learners

The **Machine Teaching** project advances AI-supported tutor creation through two complementary systems, TEACH and LEARN, designed to help instructors rapidly build, customize, and refine intelligent tutors for adult learners. These systems make it possible for teachers to author tutor interfaces, expert models, and procedural logic without needing programming expertise, enabling instruction that is aligned with real-world practice and continuously adaptable to instructor needs.

Creating Tutors Through TEACH and LEARN

TEACH: A Generative + Direct-Manipulation Interface Builder

- TEACH allows instructors to design tutor interfaces using natural-language prompts and a drag-and-drop editor. Instructors describe the tutor they want, and a large language model generates an initial interface layout. Teachers can then refine the design manually or request additional edits, making the creation process fast, iterative, and intuitive.

LEARN: Interactive Task Learning for Expert Models

- LEARN guides instructors through modeling problem-solving steps using the L.E.A.R.N. framework. Instructors validate or modify each step to help the system learn the structure of a problem. A built-in visualization shows the tutor's hierarchical task network (HTN), giving teachers insight into how the agent interprets tasks and how its knowledge is structured.

STAND: More Predictable AI Learning With Human Feedback

The team developed STAND, a new interactive task-learning approach that learns rule conditions efficiently and provides clear signals about when more instruction is needed. STAND reduces regressions, improves stability, and helps instructors know exactly when an AI agent has been fully trained on a skill — addressing a major challenge identified in earlier ALOE work.

Key Insights from Teacher Engagement

Engagement with TCSG nursing instructors revealed a critical need: every instructor teaches differently, and each requires personalized control over tutor behavior.

Two findings stand out:

- Instructors vary in how they want material taught, even in standardized domains like drug calculations.
- Individual instructors change their pedagogy across semesters, meaning static tutors cannot keep up.

These insights confirm that adoption of AI tutors depends on giving instructors the ability to personalize tools to their needs, a core motivation for TEACH, LEARN, and STAND.

Contributions to Adult Learning

The project identifies personalization needs unique to adult-education instructors and develops tools that enable:

- AI tutors aligned with real-world tasks and job-specific workflows
- Instruction that adapts to learner goals and workplace relevance
- Structured support that reduces cognitive load and improves self-paced learning

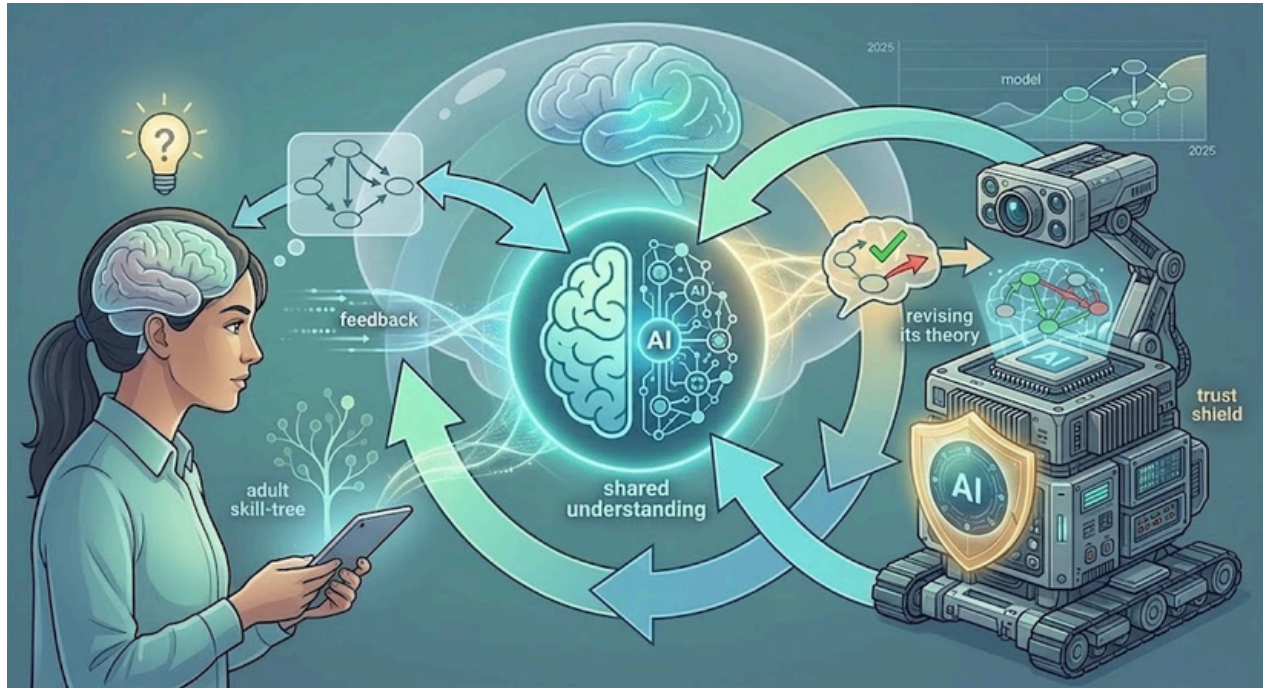
Collaboration and Team Development

The Machine Teaching effort complements the Apprentice Tutors project, with shared instructor meetings and co-design activities informing both systems. The project also supports growth for PhD students and postdocs through research, platform development, and instructional collaborations.

Mutual Theory of Mind (MToM)

Building Shared Understanding Between Humans and AI

The **Mutual Theory of Mind (MToM)** project explores how humans and AI systems can develop shared mental models of one another to improve communication and interaction. While humans naturally form beliefs about how AI systems work, current AI systems do not build models of how humans understand the AI. This limits dialogue, trust, and collaborative reasoning in AI-mediated learning environments.



A Framework for Mutual Theory of Mind

Building on foundational cognitive science, the project draws on a three-stage MToM process model developed in earlier research. This model describes how an AI system can:

1. form an initial theory of the user,
2. reveal that theory to the user through feedback, and
3. revise its theory based on the user's response.

This recursive process forms the basis of mutual understanding.

Early Computational Implementation

In 2025, the team created the first computational implementation of the MToM process model. This initial version handles a limited set of user feedback types within SAMI's social-recommendation workflow, but it establishes a foundation for expanding second-order reasoning in future AI agents. The implementation is grounded in metacognitive architectures designed to track, reflect on, and correct the AI's own reasoning.

Relevance for Adult Learning

MToM supports AI-enabled learning tools by helping systems better interpret learner intent, adjust feedback, and build transparent explanations. This work lays the foundation for more adaptive, trustworthy AI systems that can support adult learners across diverse contexts, aligning with AI-ALOE's broader goal of enhancing human-AI communication and learning.

SAMI Strengthens Social Connection in Online Adult Learning

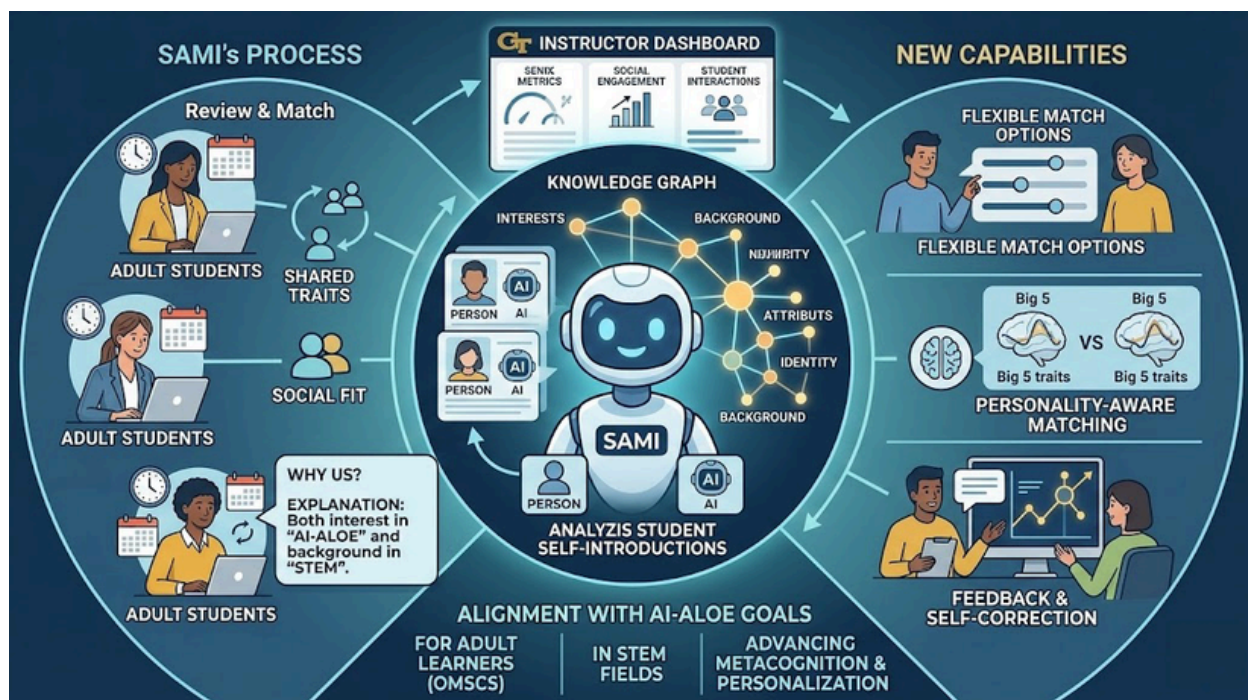
SAMI is an AI agent designed to bridge the social gaps in online learning. By analyzing student introductions via knowledge graphs, SAMI recommends classmates with shared interests and backgrounds, fostering a sense of belonging for busy adult learners.

Core Functionality

- AI-Mediated Matching: Recommends peers based on shared traits and identities.
- Metacognitive Explanations: Provides transparent reasoning for why specific matches were made.

New Capabilities

- Customizable Matching: Students can weight specific attributes (e.g., career vs. hobby) to prioritize their own match criteria.
- Personality & Self-Correction: Incorporates personality-aware insights and allows users to flag incorrect matches to refine future accuracy.
- Instructor Dashboard: Gives faculty visibility into social engagement and identifies students needing more support.
- Research-Driven Design: New studies on student preferences ensure the AI's explanations and matches align with human expectations.



SAMI supports key AI-ALOE priorities by:

- **Serving** adult learners with limited time for social engagement
- **Operating** primarily in online programs such as Georgia Tech's OMSCS
- **Supporting** STEM fields and AI-focused education
- **Advancing** work in metacognition, personalization, and Mutual Theory of Mind
- **Contributing** to research on socially aware AI systems

Supporting Adult Learners

Adult learners often de-emphasize social activity, even though connection plays an important role in motivation and persistence. By offering low-effort, meaningful matches and clear explanations, SAMI helps adult learners build community and feel more supported in online environments.

Self-Explanation

Advancing Self-Explanation for More Transparent AI

The **Self-Explanation** project focuses on improving the transparency, trust, and usability of AI systems by enabling agents to explain how they reason. The goal is to help adult learners better understand AI recommendations and to support clearer, more effective human–AI interaction.

Why Self-Explanation Matters

AI adoption often depends on users' trust in the system. Users are more likely to trust AI recommendations when the system can explain how decisions were made. Self-explanation allows an AI agent to reveal its reasoning in a concise, structured, and understandable way, helping learners decide when and how to rely on AI support.

A Metacognitive Architecture for AI Explanations

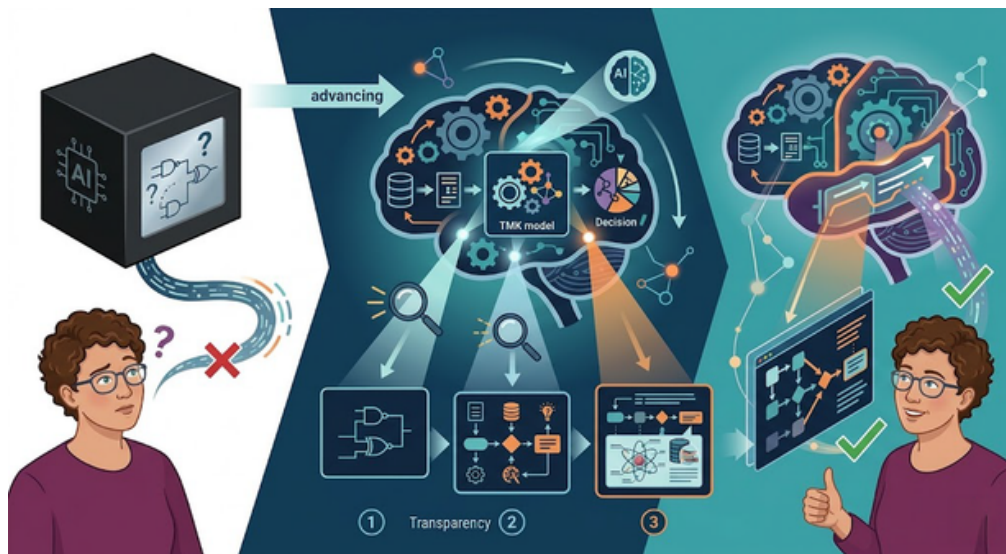
The team developed a metacognitive architecture based on Task–Method–Knowledge (TMK) models. These models allow an AI agent to represent what it knows, how it reasons, and how steps relate to underlying domain concepts. When a learner asks for an explanation, the agent uses this structure to introspect and generate a clear, accurate self-explanation.

The first implementation was introduced in SAMI, where early classroom use showed that learners responded positively to more transparent AI behavior. The team then extended the same architecture to VERA, demonstrating the model's scalability across more complex, multi-step tasks.

Research Insights From Human-AI Interaction

A series of studies explored what kinds of explanations users want and how those needs change. Findings showed:

- In everyday scenarios, many users prefer minimal explanations due to limited time and attention.
- Users' explanatory needs evolve over time, shifting from basic clarifications to deeper insight.
- The effectiveness of explanations depends on aligning properties such as conciseness or detail with the user's goal, such as persuasion, understanding, or verification.



These insights informed refinements to the architecture so that AI-generated explanations can evolve along with user needs, supporting a more adaptive, human-centered interaction style.

Toward Human–AI Co-Evolution

Based on this work, the team is advancing a model in which AI explanations and user expectations evolve together. This co-evolution helps AI systems provide the right amount of information at the right time, supporting transparency without overwhelming users, and helping adult learners stay in control of their learning experience.

SMART Advances AI-Powered Concept Learning Across Courses

The **SMART** team continues to advance AI-powered concept learning and instructional design support across nursing, STEM, and graduate education programs. Through three interconnected efforts: **AI-Empowered Instructional Design Practices**, **AI-Powered Concept Learning in Professional Training**, and **Five Years of Concept-Learning Experiments**. SMART strengthens adult learners' ability to understand complex content and supports instructors in integrating AI meaningfully into their courses.

AI-Powered Concept Learning Across Disciplines

SMART helps adult learners summarize difficult material, refine misconceptions, and deepen conceptual understanding through iterative AI feedback. The system has been deployed across nursing, physics, biology, computer science, and English courses, benefiting more than 2,200 learners over five years. In nursing programs, SMART supports core courses such as Medical Surgical I, Pathophysiology, and Advanced Pathophysiology.

Instructional Design Insights and Case Studies

One branch of SMART research documents how instructors adapt AI tools into their courses. Case studies show shifts in teaching practices such as:

- Using SMART as a pre-class or in-class concept-building activity
- Coordinating SMART with Jill Watson to support revision and question-asking
- Integrating AI into instructional design graduate courses to help learners analyze content, generate questions, and evaluate AI tools

These cases are informing a developing framework for AI-supported instructional design.

Impact on Learning Behaviors and Engagement

Across deployments, SMART has revealed distinct learner behavior patterns and different pathways adults use to improve conceptual models. Studies show:

- Iterative revision improves conceptual clarity
- SMART-supported learners outperform non-supported peers on concept-learning measures
- No demographic gaps emerge in learning outcomes

SMART's design has also evolved in response to challenges in monitoring pre-class engagement, prompting a shift toward guided in-class summarization activities in undergraduate courses.

Confluence With Jill Watson

In several nursing courses, SMART and Jill Watson worked together to support:

- Concept clarification
- Study preparation
- Summary writing and revision
- Question-asking about concepts and logistics

This confluence allowed learners to move fluidly between building conceptual models and asking targeted questions.

Research on Personalized, Equitable AI Learning

Across five years of experiments, SMART has contributed to research on:

- Approaching the two-sigma effect of personalized learning
- Tracking learner engagement patterns across disciplines
- Understanding how revision behavior reflects conceptual growth
- Supporting equitable learning outcomes for diverse adult learners

These insights strengthen the design of AI tools that are transparent, equitable, and aligned with adult learning needs.

Tech Well-Being Assessment

User-Engaged Design Team Assesses Well-Being Across AI-ALOE Technologies

The **User-Engaged Design for Human Well-Being (U-ED)** team evaluates how AI-ALOE technologies affect users and stakeholders, using the framework defined in IEEE Standard 7010-2020. The standard outlines a structured, multi-step process for assessing human well-being impacts in autonomous and intelligent systems.

Over the past four years, the U-ED team has applied this process to multiple AI-ALOE technologies. The chart below summarizes progress across projects from 2021 through October 2025.

Demographic Survey for Focus Groups

To understand how AI tools affect different learner groups, the team developed a brief, non-intrusive demographic survey. It was used first in SAMI focus groups and later in VERA sessions. Results showed:

- SAMI focus groups were largely homogeneous across most demographic criteria, with variation only in gender and racial identity.
- VERA focus groups were small but more diverse in gender and racial/ethnic identity, though less varied in age and geography.

Because the survey does not collect personally identifiable information, feedback cannot be traced back to specific demographic categories.

Supporting Diverse and Marginalized Learners

The team aims to ensure that AI-ALOE tools reflect the perspectives of users who may not otherwise be heard. Recruiting diverse participants remains a challenge due to variations in college demographics across deployment sites. To broaden representation, the team partnered with instructors at institutions in suburban, small-city, and global online contexts.

Key Impact Highlights

Two major findings emerged from the team's well-being assessments:

- Apprentice Tutor Redesign Improves Outcomes
- An initial review identified negative well-being impacts. After the team recommended changes, developers redesigned the tool and redeployed it. Follow-up interviews showed concerns had been addressed, marking a significant success for user-engaged AI design.
- SAMI Shows Strong Well-Being Benefits—and Raises Important Questions
- SAMI supports social connection for online learners by encouraging relationship-building and reducing isolation. At the same time, focus groups revealed a broader challenge for social AI: how to support engagement without replacing real opportunities for learners to develop interpersonal skills.



The project supports the professional growth of graduate researchers, including mentoring, applied research experience, and structured preparation for careers in user-centered AI, UX research, and responsible technology design. The U-ED team ensures responsible design of AI-ALOE technologies by applying a refined version of the IEEE 7010-2020 well-being assessment process. Their work has shaped the design of multiple AI-ALOE tools, identified risks and opportunities for user impact, and supported the institute's commitment to ethical, human-centered AI. More information is available at aialoe.org/ai-ethics.

VERA Expands Inquiry-Based Learning and Model-Building Skills for Adult Learners

VERA (Virtual Ecological Research Assistant) is an interactive learning environment designed to help adult learners think like scientists through inquiry-based learning. Many adults have limited exposure to scientific modeling, so VERA provides structured support for exploring models, adjusting parameters, and understanding how simulations generate insight.

Helping Adults Learn Through Inquiry

Effective inquiry in VERA requires both domain knowledge and modeling skill. To support these needs, VERA provides multiple coaching agents that offer cognitive scaffolding. These include assistance with domain facts through the Encyclopedia of Life interface and guidance for choosing simulation parameters.

Because adult learners often avoid reading manuals, VERA now includes AskJill, a question-and-answer agent trained on the quick start guide and reference manual. Learners can ask procedural or conceptual questions and receive clear, targeted support.

Introducing Self-Explanation for Greater Transparency

In a recent pilot, VERA added Self Explanation, a feature that helps the system articulate its goals, tasks, and reasoning. This capability draws on Mutual Theory of Mind and metacognition, using a structured Task-Method-Knowledge representation combined with large language models. By revealing how it makes decisions, VERA helps learners build trust and understand how scientific reasoning works.

Expanding Beyond Ecology

To broaden VERA's usefulness, the team developed VERA for Personal Financial Planning (FP). This version adapts the ecological modeling framework to financial scenarios, enabling users to build models and run simulations in a new domain. This work moves VERA toward VERA Meta, a general-purpose platform that can support conceptual modeling in any domain with a defined ontology and parameter space.

Personalized Learning Dashboards

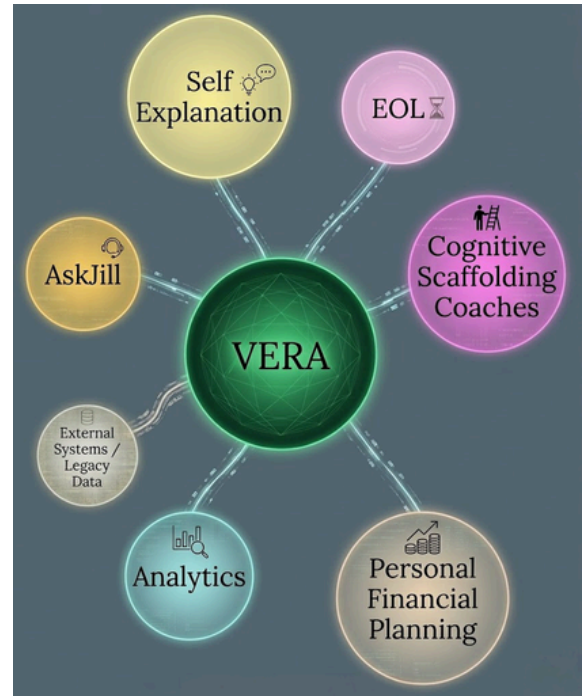
VERA is integrating with the A4L 2.0 feedback loop and developing Personalized Learning Dashboards for three audiences:

- Learners receive insights into their progress and interaction patterns compared with peers.
- Teachers can identify learners who need additional guidance and provide targeted support.
- Researchers can conduct longitudinal analyses of how AI agents work together to enhance learning outcomes.

Supporting Adult Learners' Needs

VERA focuses on helping adult learners who need flexible, guided support for scientific reasoning and modeling. By combining domain coaching, Q and A support, self-explanation, and personalized dashboards, VERA gives learners the tools to explore complex systems with confidence.

VERA has been deployed across multiple courses at Georgia Tech and NGTC, supporting online and blended learning environments. Its design aligns with AI-ALOE priorities, including support for adult learners, online learning, foundational AI, and Mutual Theory of Mind.



Visualization

VERA – personalized learning through visualization dashboards

Adult learners have a need for tools tailored to their learning pathway. **VERA (Virtual Ecological Research Assistant)** was created to meet this growing need in adult education to improve student outcomes. Through VERA, the AI-ALOE team created Personalized Learning Analytic Dashboards.

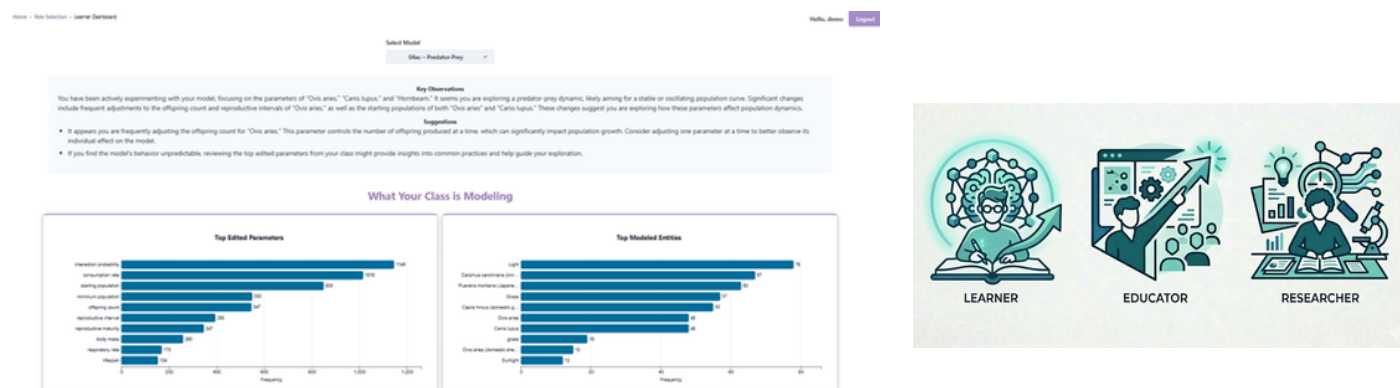
These personalized dashboards use generative AI to highlight and identify learner behavior. The observations are then shared with the student and the teacher. Both can use this information to adjust and improve the impact of the course. Adult learners can compare their results with other models – also generated by VERA. Additionally, there is visualization of data across time – a Research Dashboard.

Learner Dashboard

The learner dashboard uses generative AI to highlight how individual learners model compared to their classmates. It shows key observations about the learner's modeling behaviors and, when certain criteria are not met, generates suggestions for how to adjust or revise their model.



Teacher Dashboard



The teacher dashboard provides a retrospective view of how learners' interactions with VERA relate to class performance. It visualizes:

- Score ranges relative to interaction counts
- Grade distributions relative to model acceptability
- Modeling differences across class sections

Research Dashboard

The research dashboard supports longitudinal studies of human–AI interaction (xAI). It visualizes how learners allocate time across VERA activities in different classes and how their engagement evolves over time.



Staff Spotlight



CHARMAIN ALSTON
GRANTS ADMINISTRATOR



Charmain Alston joined the Georgia Tech in 2017 after 19 years at the National Science Foundation (NSF), where she served as a team lead and program analyst in the Computer and Information Science and Engineering Directorate. Before NSF, she worked for the U.S. Army and the U.S. Marine Corps and held assignments with the National Aeronautics and Space Administration (NASA) and the Executive Office of the President, including as an intern budget specialist in the Office of Management and Budget (OMB). Her career in federal service and research administration has connected her with leading engineering, science, and technology researchers and institutional leaders at Georgia Tech.



Q&A

What are your main responsibilities with AI-ALOE?

As the grant administrator for AI-ALOE, my primary responsibility is providing budget oversight and accountability to help ensure our principal investigators, staff, and partners remain in compliance with federal, state, and institutional policies. I manage both pre- and post-award activities and work to make sure grant actions and expenditures are handled appropriately so the project stays on solid fiscal footing.

How did you become interested in your field?

Honestly, I never imagined this as my career path. I have never been particularly drawn to working with numbers and would probably rather spend my time volunteering or touring with a theater company. However, my curiosity about the “why” and “how” behind science and technology, along with the people who help make that work possible behind the scenes, eventually led me into spaces where I get to work alongside leaders dedicated to answering those big questions.

What do you find most rewarding about work with AI-ALOE?

One of the most rewarding aspects is learning about new breakthroughs and tools in AI. I also value the opportunity to meet and work with researchers and innovators whose work I might never have encountered otherwise.

Member Spotlights



BECCA STOREY
GRADUATE RESEARCH ASSISTANT

Rebecca Storey is a Graduate Research Assistant with AI-ALOE and a doctoral candidate in psychology at Georgia Tech. Her work focuses on adult learning in online and workplace environments, with an emphasis on how individuals gain new skills or transition into new roles over the lifespan. Rebecca joined the PARK Lab in 2023 after completing her bachelor's degree in psychology with a concentration in management studies from St. Olaf College, where she conducted her first upskilling project during an HR internship.

Within AI-ALOE, she analyzes survey data to understand how confidence, learning orientation, and related factors influence adults' openness to AI-assisted learning. Her dissertation examines hybrid intelligence in the workplace, studying how collaboration between humans and AI may change the relationship between cognitive abilities and task performance.

Outside of work, Rebecca enjoys hiking with fellow psychology students, attending live music, cooking, and walking on the Atlanta Beltline.



RACHEL SIKES, MSN, RN
AI-ALOE TEACHING FELLOW

Rachel Sikes is a Registered Nurse with more than 18 years of experience in critical care and over 10 years in nursing education. She serves as an instructor and Program Coordinator for the ASN program at Southeastern Technical College. Throughout her career, she has taught Pharmacology, Foundations in Nursing, Maternity and Pediatric Nursing, and Medical-Surgical Nursing.

As an AI-ALOE Teaching Fellow, Rachel collaborates with development teams to support the effective use of AI tools in nursing courses. She has incorporated iTELL and Jill Watson into her pharmacology instruction and values the ongoing commitment to improving these tools to meet the diverse needs of students.

Course Using AI-ALOE

Fall 2025: RNSG 1018 Pharmacological Concepts and Drug Calculations


Conferences and Talks



Ashok Goel Discusses AI for Reskilling and Workforce Development at AAAI 2026



Min Kyu Kim Shares Insights on AI in Education at Seoul National University



Ashok Goel Highlights AI Innovations at UPCEA's SOLAR 25 Conference



AI-ALOE's Rahul Dass Showcases Ivy at UPCEA SOLAR 25

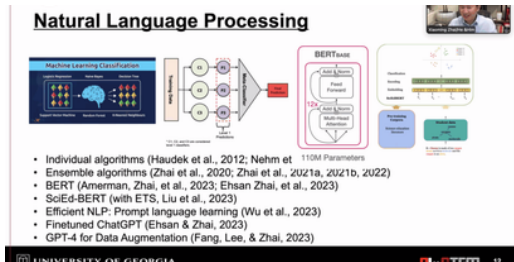


AI-ALOE Leadership at the Lifetime Learning Symposium



AI-ALOE's Sandeep Kakar Showcases Jill Watson at UPCEA SOLAR 25

Research and Innovation




[AI for Next Generation Science Education With Xiaoming Zhai](#)




[Lingqing Wang at CHI 2025](#)



[Karan Taneja Explores Multimodal Conversational AI for Education](#)





[AI-ALOE Members Present Research at ISLS 2025](#)



[Using AI in Research: Measuring Instructional Practices with Barbara Means](#)

Featured Events

AI-ALOE hosted and participated in a range of events that brought together researchers, educators, and collaborators from across the community. These gatherings created opportunities to share research, highlight ongoing projects, and strengthen connections among members and partners. Through showcases, collaborative meetings, and participation in major conferences, AI-ALOE continues to foster dialogue and advance work at the intersection of AI and learning.




[AI-ALOE Showcase Draws National Audience for Afternoon of Research and Innovation](#)



[AI-ALOE Hosts Virtual Research Showcase](#)



[AI-ALOE Hosts Virtual Mini Retreat to Advance Research and Collaboration](#)



[AI-ALOE Members attend SAIL 2025](#)

Upcoming Talks



The New Hybrid: From AI Tools to Pedagogical Systems

Speaker: Mairéad Pratschke, PhD

Independent Researcher and Advisor; Visiting Professor at the London School of Economics and Political Science (LSE) Data Science Institute; specializes in AI strategy and learning system design

Date: Thursday, May 14

Time: 11:00 AM – 12:00 PM

Location: Centergy Building, Tuff Suite, 10th Floor



AI and Lifelong Learning: Building the 60-Year Curriculum for the Fourth Industrial Revolution

Speaker: Sae Schatz

Founder, Knowledge Forge LLC; Board Member, Institute for Infrastructure and Interoperable Data in Learning (I2IDL)

Date: Thursday, May 14

Time: 2:00 PM (Talk) | 3:00 PM (Reception)

Location: CODA Building, 9th Floor Atrium




AI-ALOE invites you to join our new community group on LinkedIn!

JOIN NOW

Publications

- Bae, Y., Kim, J., & Kim, M. (2026). Four years of AI-powered formative assessment: Exploring the 2-sigma effect and fairness. *Proceedings of the International Conference of the Learning Sciences (ICLS 2026)*. (Accepted)
- Cao, L., Etemadi, A., Wheeler, M., & Dede, C. (2025). Keeping the glass box transparent: Comparing expert and AI-generated ratings and feedback in stealth assessment for judgement-focused negotiation simulations. *Journal of Research on Technology in Education*, 57, 1–23. <https://doi.org/10.1080/15391523.2025.2568526>
- Dede, C. (2024). Guest editorial on intelligence augmentation: The owl of Athena. *IEEE Transactions on Learning Technologies*, 17, 2154–2155. <https://ieeexplore.ieee.org/document/10726640>
- Dede, C. (2026). What works when for whom under what conditions: Learning engineering as an enabler of component-based research. In S. D. Craig et al. (Eds.), *Proceedings of LERN 2026*. <https://doi.org/10.59668/2551.25412>
- Etemadi, A., Tassis, A., Wheeler, M., & Dede, C. (2025). Illustrative AI-based resources for learning to negotiate. Next Level Lab, Harvard Graduate School of Education. <https://nextlevellab.gse.harvard.edu/learning-modules/negotiation-ai-project/>
- Gupta, A., Reddig, J., Calò, T., Weitekamp, D., & MacLellan, C. J. (2026). Beyond final answers: Evaluating large language models for math tutoring. In *Proceedings of the International Conference on Artificial Intelligence in Education (AIED 2026)*. https://doi.org/10.1007/978-3-031-98414-3_23
- Haddadian, G., Han, H., Kim, J., Abdeen, M. S., & Kim, M. K. (2025). Exploring AI-generated expert models: Instructor interaction and learner perceptions in a physics class. In A. Rajala et al. (Eds.), *Proceedings of the 19th International Conference of the Learning Sciences (ICLS 2025)* (pp. 1684–1688). International Society of the Learning Sciences. <https://doi.org/10.22318/icls2025.213524>
- Haddadian, G., Kim, M. K., & Haddadian, N. (2025). A systematic review of automated writing evaluation tools in argumentative writing for English as a foreign language education. *Research Synthesis in Applied Linguistics*, 1–51. <https://doi.org/10.1080/29984475.2025.2598266>
- Holmes, L., Morris, W., Crossley, S. A., & Choi, J. S. (2026). Assessing fairness in finetuned scoring models with demographically restricted training data. *Assessing Writing*. (Accepted)
- Im, S., Kim, M. K., & Heo, J. (2025). Digital literacy in physical education: A systematic review of frameworks and development areas. In A. Rajala et al. (Eds.), *Proceedings of the 19th International Conference of the Learning Sciences (ICLS 2025)* (pp. 2883–2885).
- Kanfer, R., Lyndgaard, S. F., & Storey, R. (2025). Critical considerations in human-technology partnerships for lifelong learning. In T. S. Behrend (Ed.), *Human-technology partnerships at work*. Cambridge University Press.
- Kim, J., Galindo-Parra, M. F., Kim, S., Bae, Y., Hendry, T. M., Galen, S. S., & Kim, M. K. (2025). Design-based research for scenario-based, generative AI-augmented simulation in nursing education. In A. Rajala et al. (Eds.), *Proceedings of the 19th International Conference of the Learning Sciences (ICLS 2025)*. <https://doi.org/10.22318/icls2025.571004>
- Kim, J., Islam, M., & Goel, A. (2026). A metacognitive architecture for correcting LLM errors in AI agents. *Proceedings of the AAAI Conference on Artificial Intelligence (Innovative Applications of AI)*.
- Kim, J., Islam, M., & Goel, A. (2026). A metacognitive architecture for theory-of-mind revision in AI agents. *AAAI 2026 Theory of Mind Workshop*.
- Kim, S., Han, H., & Kim, M. (2026). Design principles for AI-supported chatbots to scaffold problem-based learning in undergraduate nursing education. *Proceedings of the International Conference of the Learning Sciences (ICLS 2026)*. (Accepted)

Continue to next page 

Publications Cont.

- Lee, J. C.-K., Dede, C., Wang, M., & Li, X. (2025). Building trust in AI through dialogues with Chinese ethics: Toward ethical partnerships in education. *IEEE Transactions on Learning Technologies*, 18, 833–841. <https://doi.org/10.1109/TLT.2025.3604616>
- MacLellan, C. J. (2025). Model human learners: Computational models to guide instructional design. *Proceedings of the Cognitive Science Society (CogSci 2025)*. <https://doi.org/10.48550/arXiv.2502.02456>
- Morris, W., Holmes, L., & Crossley, S. (in press). Psychometric analysis of LLM differential sensitivity to word predictability using explanatory MFRM. *Proceedings of the Psychometrics Workshop at LAK*.
- Noroozi, O., Haddadian, G., Gao, X., Schunn, C., Alqassab, M., & Banihashem, S. K. (2025). The value of GenAI for peer feedback provision: Student perceptions and impacts. *International Journal of Educational Technology in Higher Education*, 22(61). <https://doi.org/10.1186/s41239-025-00558-6>
- Rahimi, S., Babae, M., Esmailigoujar, S., & Dede, C. (in press). Generative artificial intelligence and creativity in K-12 education: A scoping review. In R. Beghetto (Ed.), *Oxford handbook of human creativity and generative AI*.
- Rahimi, S., Dede, C., Esmailigoujar, S., & Babae, M. (2025). Augmenting human creativity with responsible and ethical use of generative AI. In M. Worwood & J. C. Kaufman (Eds.), *Generative artificial intelligence and creativity: Possibilities, precautions, and perspectives*. Academic Press.
- Reddig, J., Arora, A., & MacLellan, C. J. (2025). Generating in-context, personalized feedback for intelligent tutors with large language models. *International Journal of Artificial Intelligence in Education*. <https://doi.org/10.1007/s40593-025-00505-6>
- Reddig, J., Moon, S., Crutcher, K., & MacLellan, C. J. (2026). AI unplugged: Embodied interactions for AI literacy in higher education. *Proceedings of the AAAI Conference on Educational Advances in Artificial Intelligence (EAAI 2026)*. <https://tail.cc.gatech.edu/files/reddig-eaai-2026.pdf>
- Shen, S., Wise, A. F., Morris, W., Holmes, L., Hinze, S. R., & Crossley, S. (in press). The effects of AI feedback on college students' reading and writing performance in an intelligent text framework. *Proceedings of the 16th International Conference on Learning Analytics & Knowledge (LAK)*.
- Smith, G., & MacLellan, C. J. (2025). L.E.A.R.N: A hybrid architecture for language-guided induction of hierarchical task networks. *Proceedings of Advances in Cognitive Systems 2025*. <https://openreview.net/pdf?id=aKJl1p8hWb>
- Thajchayapong, P., Carbonaro, S., Couper, T., Helmick, B., Rugaber, S., & Goel, A. (in press). Evolution of A4L: A data architecture for AI-augmented learning. In W. Lawless et al. (Eds.), *Generative AI risks and benefits within human-machine teams*. Elsevier Press. <http://arxiv.org/abs/2511.11877>
- Wang, L., Gao, Y., Anyi, C. L., & Goel, A. (2026). Futuring social assemblages: How enmeshing AIs into social life challenges the individual and the interpersonal. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems (CHI 2026)*. <https://doi.org/10.1145/3772318.3790279>
- Weitekamp, D., MacLellan, C. J., Harpstead, E., & Koedinger, K. (2025). Decomposed inductive procedure learning: Learning academic tasks with human-like data efficiency. *Proceedings of the Cognitive Science Society (CogSci 2025)*. <https://arxiv.org/abs/2505.10422>
- Weitekamp, D., Siddiqui, M., & MacLellan, C. J. (2026). TutorGym: A testbed for evaluating AI agents as tutors and students. *Proceedings of the International Conference on Artificial Intelligence in Education (AIED 2026)*. https://doi.org/10.1007/978-3-031-98420-4_26
- Zappatore, M., Wang, M., Dede, C., & Li, X. (2025). Intelligence augmentation and the future of education. *IEEE Transactions on Learning Technologies*, 18. <https://doi.org/10.1109/TLT.2025.3639317>

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



UNC GREENSBORO



VANDERBILT UNIVERSITY



The material is based upon work supported by the National Science Foundation under Grant No. 2247790 and Grant No. 2112532. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.