iSAT Snapshot

NSF AI INSTITUTE FOR STUDENT-AI TEAMING Summer 2025

In Brief . . .

A quick look at our activities this quarter!

1. NSF iSAT co-hosted a day-long Multimodal Multiparty Learning Analytics (MMLA) workshop with partial support from Google.org at the Educational Data Mining Conference in Palermo, Sicily.
2. Strand 3 shared new research and curriculum resources at a range of conferences: internationally at the International Conference of the Learning Sciences (ICLS) and the International Conference on Computer-Supported Collaborative Learning (CSCL), nationally at the Computer Science Teacher Association (CSTA) annual conference, and locally at the Rocky Mountain Computer Science (RockCS) Conference.
3. The paper "From Discourse to Dynamics: Understanding Team Interactions Through Temporally Sensitive NLP" authored by Seehee Park, Danielle Shariff, Mohammad Amin Samadi, and Nia Nixon along with iSAT-PI Sidney D'Mello received the Best Short Paper Award at the International Conference on Educational Data Mining (EDM 2025).
4. The paper "Towards a robust automated system of detecting collaborative problem solving markers in a small group collaborative setting" by Videep Venkatesha, Mariah Bradford and Nathaniel Blanchard was included in the "Late Breaking Results – Best Papers" Nominees at the 26th International Conference on Artificial Intelligence in Education (AIED).
5. Sandra Sawaya alongside iSAT-PI Sidney D'Mello and other affiliated researchers were finalists for two best paper awards at the 26th International Conference on Artificial Intelligence in Education (AIED).

iSAT-PI Sidney D'Mello presents at the 2025 Educational Data Mining Conference in Palermo, Sicily.

From the PI

It's hard to imagine that iSAT in Phase I (2020-2025) is coming to an end. To celebrate our work, we organized a three-hour public-facing showcase attended by several NSF personnel in addition to the site visit team (SVT), who evaluates iSAT each summer. We presented our key accomplishments, including the development, testing, and fielding of two innovative AI Partners (the Community Builder and the Jigsaw Interactive Agent) in classrooms in close coordination with teachers. We explained how embedding the partners in novel AI literacy curricula helped develop new capacities for AI innovation for more than 6,000 middle-school students and their teachers. Lastly, we discussed our research contributions including advances in AI for real-world classrooms, developing the new science of student-AI teaming, and pioneering new models of teacher professional learning. Based on the feedback received, the NSF and SVT were very impressed with our progress and expressed considerable enthusiasm for this work to continue.

This leads us to Phase II of iSAT (2025-2030), which we launched in September. In this next phase, iSAT will address the urgent national need to develop a next-generation AI workforce while simultaneously amplifying our impact as a nexus by connecting research, education systems, and industry to accelerate innovation and workforce readiness in AI. Phase II provides us with an opportunity to reimagine our research agendas in the brave new world of GenAI. It encourages us to review and revive our commitments to our students/educators, to each other, and to our scientific disciplines. It enables us to streamline our processes, enhance the quality of our work, and help members of our team reach new potentials. We will be discussing these and related issues at an in-person retreat in September, where we aim to develop the outlines of an impactful strategic and implementation plan to guide the institute in Phase II.

Feature: Workshop on Multimodal Multiparty Learning Analytics

An iSAT-led Cross-Institute Collaboration at Educational Data Mining (EDM) 2025

Multimodal, Multi-party Learning Analytics (MMLA)–combining multiple data streams from several participants–has great promise in providing richer insights into learning processes. Yet there are significant challenges in using MMLA in research and in deploying it in educational settings. Further, while there are a number of organizations individually working on MMLA approaches, the field is still young and researchers can benefit from sharing and from the cross-pollination of methods. To address these challenges, NSF iSAT, in collaboration with the NSF AI Institute for Engaged Learning (EngageAI), hosted a day-long Multimodal Multiparty Learning Analytics (MMLA) workshop at the Educational Data Mining (EDM) Conference in Palermo, Sicily. The workshop was partially sponsored with support from Google through NSF's AI Institutes Virtual Organization (AIVO).

Bringing Researchers Together

The objectives of the workshop were to bring together a diverse group of researchers who are researching and developing MMLA for education, to share their methodology, and to illustrate techniques on their own data. A further objective was to foster enhanced collaboration across NSF AI for Education Institutes. Presenters submitted either a research paper (a description of a particular analytic method), empirical results, a data set, or a position paper which argued for a particular position of need in the field. The workshop consisted of nine presentations spanning five NSF AI for Education Institutes including EngageAI, INVITE, AI-ALOE, AI4ExceptionalEd, and iSAT, as well as other research organizations working in the area. As an example, Ekta Sood (iSAT) presented work done collaboratively with Joyce Fontalese (EngageAI) and Utkash Kumar (AI4ExceptionalEd) on a proof-of-concept for building lightweight multimodal encoders through cross-institutional efforts. The format of the workshop incorporated short-presentations followed by extended time for in-depth discussions with each presenter. This allowed greater knowledge sharing, building common ground, and identifying potential collaborations. In addition, a panel session was held with participants from four NSF AI for Education Institutes (EngageAI, iSAT, INVITE, AI-ALOE) where the panelists discussed their research in MMLA and how it related to their institute's broader goals.

Charting a Course to Advance the Field

Overall, the workshop fostered greater knowledge-sharing across the AI Institutes as well as with other researchers in AI in education. It addressed key considerations needed for the field to continue to advance, including data collection of different modalities, techniques for integrating multiple data sources from multiple individuals, advances in AI-based analysis techniques, multimodal analytic pipelines, and approaches to addressing data privacy and ethical frameworks. Participants further discussed approaches to addressing challenges in the field, including multimodal fusion techniques, validating performance of models of complex educational interactions, and scaling and implementation of these approaches in the wild, while respecting the privacy constraints of classroom data. Follow-up discussions based on the workshop will be held during the AI4Ed workshop at the NSF SAIL meeting in October.

Presenters and panelists at the 2025 Workshop on Multimodal Multiparty Analytics at EDM 2025 in Palermo, Sicily.

Feature: Exploring Generative AI in Education

Insights from a Workshop with Educators

The rapid emergence of Generative AI (GenAI) technologies, such as ChatGPT, has created new and urgent challenges for schools. These tools have entered classrooms with little warning, raising pedagogical, ethical, and institutional questions well before educators and policymakers have had the opportunity to respond. According to Reich and Dukes, GenAI is an "arrival technology"—one that appeared suddenly, disrupting the usual processes of planning and preparation that accompany the introduction of new educational tools.

Exploring GenAI as an Object of Joint Inquiry

This spring, NSF iSAT organized a four-session workshop series designed to address this very challenge. We analyzed the data from this workshop over the summer. The series brought together seven teachers from three school districts who had previously implemented our instructional materials. The aim was to explore how GenAI might serve not only as a teaching tool, but as an object of joint inquiry, something that students and teachers could investigate together as part of developing AI literacy.

The workshop had two central goals: first, to foster a community of practice among educators and researchers interested in deepening their collective understanding of GenAI; and second, to identify learning goals and classroom contexts that support students in engaging with the technology critically and creatively. Over the course of the four sessions, participants reflected on their experiences with GenAI, developed shared understandings of how these technologies work, and analyzed student perspectives on its implications. Taken together, these three topics offered meaningful entry points for engaging in critical inquiry about GenAI. They align closely with emerging research that emphasizes the importance of AI literacy as a multidimensional skill set. What emerged was a promising foundation for designing curriculum units based around students using AI in classrooms.

A Closer Look at Each Session

The first session invited teachers to share their experiences with the arrival of GenAI in their schools. These discussions revealed a mix of experimentation, caution, and curiosity. Many teachers expressed a sense of being underprepared, acknowledging that the speed of AI's adoption had outpaced institutional efforts to establish guidelines or provide professional development, affirming the need for sustained dialogue among educators and researchers as they adapt to this new technological landscape.

The second session focused on building a shared understanding of how GenAI functions. Teachers emphasized the importance of helping students understand that AI tools are not magical or autonomous, but rather data-driven systems shaped by human design choices. Participants noted the importance of equipping students with the critical thinking skills necessary to use and evaluate AI. In the third session, attention shifted toward identifying instructional opportunities. Teachers generated several candidate topics that would be both intellectually substantive and relevant to students' lives. The aim was to center inquiry-based learning around topics that encourage analysis, debate, and reflection, marking an important step toward a potential storyline unit anchored in students' own questions and concerns about AI.

The final session incorporated student voices and what they found compelling when thinking about GenAI. Three topics in particular stood out: the intersection of AI and art, the role of AI in teaching, and the environmental costs of AI development. The topic of AI and art prompted reflections around ownership, credit, and creativity. For AI and teaching, students wondered whether AI could replace teachers, what would be lost in terms of human connection, and how the role of the classroom might evolve. Students also voiced interest in the environmental implications of AI, expressing concern about the energy required to run large-scale AI systems and curiosity about how these systems work.

Looking to the Future

Scholars increasingly argue that students should not only learn how to use AI, but also how to understand its design, evaluate its outputs, and consider its societal impacts.

The NSF iSAT workshop series contributed to this growing body of work by modeling how professional learning communities can support thoughtful integration of new technologies in education. By combining educator reflection, collaborative inquiry, and student-centered data analysis, the series provided a generative framework for exploring how AI literacy might be taught in schools. It also demonstrated the value of treating teachers as co-designers by bringing their classroom expertise into conversation with research that informs the development of responsive and relevant instructional materials.

As questions surrounding GenAI continue to shape students' educational experiences and future opportunities, these types of workshops will play a crucial role in guiding responsible and informed integration. Supporting educators in navigating arrival technologies is not only about technical fluency; it is about fostering thoughtful dialogue, reasoning, and pedagogical innovation. The work of the educators who participated in this series represents an important step toward that goal.

Checking in on NSF iSAT's Research Strands

Summer Updates

Strand 1 - Listen, Understand, and Connect

This summer, the Strand 1 team focused on improvements to the Jigsaw Interactive Agent's (JIA) content-based response capabilities. Here, our primary attention was on advancing the retrieval component of our curriculum-based Retrieval-Augmented Generation (RAG) response system. We incorporated additional documents that relate to questions one and two on the Sensor Immersion (SI) task to the curriculum database. We further added a filter to remove questions from the curriculum database to improve the quality of the retrieved knowledge. Looking forward and in coordination with the Strand 3 research team, we are working on creating high quality annotations for better identification of content-based response opportunities. Specifically, we are looking for dialog spans that correspond to opportunities to provide knowledge-informed responses to students.

The speech team explored the utility for speaker diarization to utilize textual rather than auditory embeddings. The idea is that certain speakers may have a linguistic style that distinguishes them from others. Initial experiments on the SI dataset based on the StyleDistance text embedding model indicate that text embeddings do carry some information about speaker identity but less than auditory embeddings. Moreover, a combined approach yields only minimal benefit. Continuing our group's work on multimodal LLMs, we found that (a) the same trends regarding the benefits of multimodality for speech recognition that we identified in Transformers are also exhibited by Mamba, and (b) Transformer-based LLMs have more difficulty processing 2-d representations (e.g., a text raster showing what a person is talking about) than 1-d representations of the same information (e.g., a list of concatenated strings of the same content). This suggests that further research in positional encodings or attention mechanisms might be helpful to enable multimodal LLMs to process image and video content. Last, we investigated the benefits of user-provided real-time feedback on diarization and speech recognition to speech processing systems. In experiments conducted on the "AMI Meeting" dataset, we found that this feedback can deliver meaningful benefits to speech processing accuracy while requiring only very sparse feedback from the user.

The team that focuses on situated grounding worked on the following projects: (a) Annotated Dataset of Sensor Immersion Interactions; (b) Detecting Co-Attention from RGB Video; (c) Object and Action Detection for Integration with AI Partner Demo; and other nonverbal behavioral cues of engagement. This was performed in the context of work on a general multimodal behavior recognition (MMBR) system. We developed a 2D multimodal perception system comprising three main components: (1) human body (rig) detection, including posture, gesture (e.g., nodding, pointing), gaze estimation, and facial expression recognition; (2) object detection, including items such as laptops and phones, along with their movement tracking—which can be used to infer actions (and generate action annotations); and (3) human–object interaction (HOI) detection, covering behaviors such as typing, gaze target prediction, and pointing target prediction. Additionally, the system supports the implementation of custom-coded modules for detecting specific gestures as needed.

Strand 2 - Orchestrating Interactions

This summer, the Strand 2 team has been continuing to analyze the rich, multimodal data from our JIA Lab studies with more than a hundred middle and high school students. Participants worked in dyads and triads on a Sensor Immersion jigsaw activity while the Jigsaw Interactive Agent (JIA)—played by a human Wizard-of-Oz (WoZ) and by our automated JIA LLM-agent—supported the group activity. Findings from this study are currently making their way into two papers that explore the multimodal nature of collaboration in small groups. One paper examines three WoZ sessions and three Human-in-the-Loop (HITL) sessions to analyze students' verbal and nonverbal interactions using the Collaborative Problem-Solving (CPS) framework (verbal) and the Nonverbal Interactions in Collaborative-Learning Environments (NICE) framework (nonverbal), respectively. We are examining correlations between multiple modalities to better understand how students engage with one another during collaboration, with support from an AI partner. The second paper is an investigation of the eyetracking data, looking at moments of shared attention and how these moments are associated with performance outcomes (each jigsaw electronic worksheet is scored using a grading rubric) and discourse (verbosity, collaborative problem solving, and turn-taking) in the group.

Strand 3 - Engaging Youth

This summer, Strand 3 shared new research and curriculum resources on the Self-Driving Cars and AI unit, the Sensor Immersion unit, and related curriculum routines at a range of conferences—internationally at the International Conference of the Learning Sciences (ICLS) and the International Conference on Computer-Supported Collaborative Learning (CSCL), nationally at the Computer Science Teacher Association (CSTA) annual conference, and locally at the Rocky Mountain Computer Science (RockCS) Conference. A highlight was the public release of the Self-Driving Cars and AI unit at CSTA, where teachers and administrators engaged in hands-on learning about how to implement this free resource in their schools and districts. These presentations reached diverse audiences, building momentum for integrating AI-infused, inquiry-based learning into classrooms.

In parallel, we advanced several lines of analysis of classroom data. We completed an initial study of the Self-Driving Cars co-design and implementation process for a paper submitted to Special Interest Group Computer Science Education (SIGCSE) and began analyzing Student Experience Exit Ticket (SEET) data from its classroom rollout for an upcoming ISLS submission. For the Sensor Immersion unit, we are progressing with the analysis of curriculum routines during the CoBi implementation, developing comparative instrumental case studies to complement quantitative findings with the Moment of Support Analysis in Collaboration (MOSAIC) framework on teacher enactment. Together, these studies will provide a robust, multi-method understanding of how to design and implement AI-integrated curricula in ways that are both pedagogically sound and adaptable to different kinds of educational contexts.

We also conducted five separate analyses of implementation of the Online Moderation Unit and submitted these as a symposium proposal to the American Educational Research Association (AERA). One focused on how implementation research informed design changes to improve student understanding of how AI systems are trained to help moderate online communities. A second analysis describes how teachers across disciplines adapted the unit to fit their instructional goals in the professional learning (PL) sessions. A third analysis examined students' interactions with feedback from JIA, examining how students relayed and reacted to the feedback in order to explore the partner's influence on collaboration and task completion. A fourth analysis focused on nonverbal behaviors during collaborative activities to broaden the conceptualization of collaboration and inform future refinements of the AI-partner embedded units. A fifth focused on how students' performance on a learning assessment varied across classrooms. Together, these analyses provide a rich, multi-method picture of how AI-infused curricula can be implemented effectively to deepen understanding of AI among diverse learners. In addition, the findings offer policymakers, practitioners, and the public a clear view of how implementation research can guide the development of AI-infused learning in the future. In the months ahead, we will be picking up the threads of co-design, building a new unit on Generative AI, supporting usability studies of our AI partners, and continuing analyses of implementation and student outcomes.

Institute-Wide - Data Collection, Architecture & Implementation

Over the summer, the institute-wide team has been focused on analyzing classroom data, with several key areas of activity. One major line of work involves coding student group interactions for Collaborative Problem Solving (CPS) indicators and scoring student surveys that capture student experiences and perceptions related to collaboration, curricula, and use of our AI tools.

Another focus has been on the development and application of a new coding scheme to analyze teacher-led "Revisitings" that involve the Community Builder AI Partner (CoBi). This is the phase where teachers bring CoBi's AI-generated feedback into whole-class discussions. The coding is focused on how revisitings are orchestrated, what revisiting routines are used, how student-driven the conversations turn out to be, and whether students appear to be using AI partner feedback to verbalize their reflections, among other areas of interest. This work is distinct from the CPS coding and is helping the team understand how teachers integrated AI tools into their instructional practices. In a related but separate area of inquiry, the team has begun exploring the question: What is the relationship between AI partner use and classroom culture? This effort is still in its early stages, with team members reviewing pre-existing coding schemes to determine their fit for this type of analysis.

Taken together, these projects help us to better understand how students and teachers are engaging with collaborative learning and the AI-enhanced curricula, and how curriculum design and classroom routines influence that engagement.

Chelsea Chandler and Sidney D'Mello showcase iSAT's work at EDM 2025 in Palermo, Sicily.

iSAT Publications

Published / Accepted

Anderson, E., Lin, G., Farid, A., Fenech, M., Hanks, B., Klopfer, E., Doherty, E., Hirshfield, L., Ko, M., Foltz, P., Nguyen, H.,Nguyen, V., Ludovise, S., Santagata, R., Cao, L., Scardamalia, M., Soliman, D., Resendes, M., Khanlari, A., Costa, S., Chan,C.K.K., Nguyen, A. (2025). Exploring GenAI Technologies within Collaborative Learning. In J. Oshima, B. Chen, F. Vogel & S. Jarvela, Eds. International Conference on Computer-Supported Collaborative Learning (CSCL).

Anindho, S., Venkatesha, V. & Blanchard, N. (2025). A Methodological Framework for Capturing Cognitive-Affective States in Collaborative Learning. Paper presented at the EDM Workshop on Multimodal, Multiparty Learning Analytics, Palermo, IT.

Bradford, M., Krishnaswamy, N., & Blanchard, N. (2025, July). The Impact of Background Speech on Interruption Detection in Collaborative Groups. In International Conference on Artificial Intelligence in Education (pp. 118-131). Cham: Springer Nature Switzerland.

Bradford, M., Khebour, I., VanderHoeven, H., Venkatesha, V., Blanchard, N., & Krishnaswamy, N. (2025, May). Tracking Individual Beliefs in Co-situated Groups Using Multimodal Input. In International Conference on Human-Computer Interaction (pp. 139-158). Cham: Springer Nature Switzerland.

Breideband, T., Bush, J.B., Reitman, J.G., Rose, S., Weatherley, J., Penuel, W.R., D'Mello, S.K. (in press). A Feasibility and Implementation Integrity Study of the Community Builder (CoBi): An AI-based Collaboration Support System in K-12 Classrooms. International Journal of Artificial Intelligence in Education.

Breideband, T., Badiger, C., Bush, J.B., Rathod, S., Singh, T., Weatherley, J., and D'Mello, S.K. (2025). The Community Builder (CoBi) for Helping Students Practice Collaboration Skills in the Classroom. Interactive presentation at the 2025 AI in Education Conference (AIED). Palermo, IT.

Cai, J., King, B., Cameron, P., Windisch Brown, S., Eckert, M., Srinivas, D., Baker, G.A., Everson, K., Palmer, M., Martin, J., Flanigan, J. (2025). In Search of the Lost Arch in Dialogue: A Dependency Dialogue Acts Corpus for Multi-Party Dialogues. In Findings of the Association for Computational Linguistics: ACL 2025, pages 20135–20149, Vienna, Austria. Association for Computational Linguistics.

Cai, J., Wright-Bettner, K., Zhao, Z., Ahmed, S.R., Ramachandran, A.T., Flanigan, J., Palmer, M., Martin, J. (2025). LiDARR: Linking Document AMRs with Referents Resolvers. In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations), pages 426–435, Vienna, Austria. Association for Computational Linguistics.

Chandler, C. & D'Mello, S.K. (2025). Beyond Benchmarks: Toward Multimodal, Multiparty Learning Analytics for Real-World Impacts. Paper presented at the EDM Workshop on Multimodal, Multiparty Learning Analytics. Palermo, IT.

Chang, M.A., Penuel, W.R., & Philip, T.M. (2025). "It goes both ways": Envisioning mutually compassionate relations between teachers and students. In A. Rajala, A. Cortez, R. Hofmann, A. Jornet, H. Lotz-Sisitka, & L. Markauskaite (Eds.), Proceedings of the 19th International Conference of the Learning Sciences (pp. 547–555). International Society of the Learning Sciences.

D'Mello, S.K. (2025). How Theory Can Guide the Development of Multimodal, Multiparty Collaborative Learning Analytics: A Case Study of the Generalized Competency Model (GCM) of Collaborative Problem Solving. Paper presented at the Workshop on Multimodal, Multiparty Learning Analytics, Palermo, IT, July.

Doherty, E., Perkoff, E.M., von Bayern, S., Zhang, R., Dey, I., Bodzianowski, M., Puntambekar, S., Hirshfield, L. (2025). Piecing Together Teamwork: A Responsible Approach to an LLM-based Educational Jigsaw Agent. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 19, 1–17.

Foltz, P. W., Biswas, G., D'Mello, S., Sood, E., & Ashwin, T. S. (2025). Interactive Workshop: Multimodal, Multiparty Learning Analytics (MMLA). Proceedings of Educational Data Mining (EDM 2025).

Foltz, P. W., Chandler, C. & Ko, M. (2025). Supporting Productive Uncertainty in Multiparty Jigsaw Activities. Paper presented at the EDM Workshop on Multimodal, Multiparty Learning Analytics, Palermo, IT, July.

French, D., Moulder, R., Ezema, K., Wense, K. v. d., & DMello, S. K. (2025). Linguistic Alignment Predicts Learning in Small Group Tutoring Sessions. Findings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP 2025), Suzhou, China.

Khebour, I., Jung, C., Fitzgerald, J., Jamil, H., & Krishnaswamy, N. (2025). Feature Contributions to Multimodal Interpretation of Common Ground. In International Conference on Human-Computer Interaction (pp. 265-281). Cham: Springer Nature Switzerland.

Ko, M., Chandler, C., Weatherley, J., Raju, R.,, Muddaluru, R., and Foltz, P.. W. (2025). The Jigsaw Interactive Agent (JIA) for Facilitating Productive Uncertainty in Small Group Jigsaw Activities. Interactive presentation at the AI in Education (AIED2025) Conference, Palermo, IT, July.

Park, S., Shariff, D., Amin, M., Nixon, N., & D'Mello, S. K. (2025). From Discourse to Dynamics: Understanding Team Interactions Through Temporally Sensitive NLP. Proceedings of the International Conference on Educational Data Mining (EDM 2025), Palermo, Italy.

Penuel, W. R., Philip, T. M., Chang, M. A., Sumner, T., & D'Mello, S. K. (accepted for publication). Responsible innovation in designing AI for education: Shifting from personalization to collaborative problem-solving. Educational Researcher.

Sood, E., Dhar, P., Troiano, E., Southwell, R., & D'Mello, S. K. (2025). ScanEZ: Integrating Cognitive Models with Self-Supervised Learning for Spatiotemporal Scanpath Prediction. In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025). Association for Computational Linguistics.

Sood, E., Fonteles, J. H. & Kumar, U. (2025). Lightweight Encoders for Multimodal Engagement Prediction in Classrooms. Paper presented at the EDM Workshop on Multimodal, Multiparty Learning Analytics, Palermo, IT, July.

Zhu, Y., Lai, K., Khebour, I., Verhagen, M., Krishnaswamy, N., & Pustejovsky, J. (2025, May). Multimodal Situational Awareness: Neuro-Symbolic AI for Real-Time HCI in the Classroom. In International Conference on Human-Computer Interaction (pp. 454-464). Cham: Springer Nature Switzerland.

Submitted / Under Review / In Revision

D'Mello, S. K., Reitman, J., Beck-White, Q., & Andrews-Todd, J. (in revision). Integrating Frameworks of Collaborative Problem Solving.

Guan, Y., Trinh, V.A., Voleti, V., and Whitehill, J. "MLLM-based Speech Recognition: When and How is Multimodality Beneficial?" Submitted to IEEE Transactions on Multimedia.

He, X., and Whitehill, J. "Survey of End-to-End Multi-Speaker Automatic Speech Recognition for Monaural Audio". Submitted to Computer Speech & Language.

Zhang, R., Cao, J., Dey, I., Foltz, P., Palmer, M., Biddy, Q., Doherety, E., Bodzianowski, M., Tissenbaum, M., Hirshfield, L.. "JIA Fueled My Ideas": Designing An Interactive AI Partner for Assisting Small Group Collaborations among Students Aged 12-17. Submitted to International Journal of Computer Supported Collaborative Learning.