2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC
Host Sponsor UNC Charlotte
Paper ID #51823

Designing Role-Based AI Chatbot Interactions to Support Deliberative Reasoning in Engineering and Computing Education

Dr. Raymond L. Smith III, East Carolina University

Dr. Raymond L. Smith III, P.E., CPEM, PMP, is an Associate Professor of Engineering at East Carolina University, where he also serves as Director of Systems Engineering for the Center for Sustainable Energy and Environmental Engineering. He holds a Ph.D. in Industrial and Systems Engineering from North Carolina State University. Prior to academia, he held engineering and leadership roles in industry, including enterprise transformation initiatives at Lenovo and IBM. His teaching and research interests include engineering management, decision analysis, systems modeling, and engineering education. Dr. Smith is an American Society for Engineering Management (ASEM) Fellow and is active in the American Society for Engineering Education (ASEE) and the Institute of Industrial and Systems Engineers (IISE) through leadership and scholarly activities.

Dr. Ricky T Castles, East Carolina University

Dr. Ricky Castles is a professor in the Department of Engineering and the Associate Dean for Student Affairs in the College of Engineering and Technology at East Carolina University. He leads outreach and recruiting efforts for the College of Engineering and is Executive Director of the Student Success Center. His research interests are in the areas of engineering education-specifically broadening access to engineering for low-income students and supporting transfer student success, and wireless sensor networks.

Rachel Lee Rosenberg, East Carolina University

Rachel Rosenberg is an undergraduate Honors College student at East Carolina University pursuing a double major in Mathematics and Engineering with a focus in Industrial and Systems Engineering. Beyond her coursework, she is passionate about STEM outreach and finding new ways to help students engage with applied engineering and mathematics topics in the classroom.

Designing Role-Based AI Chatbot Interactions to Support Deliberative Reasoning in Engineering and Computing Education

Abstract

Large Language Models (LLMs) are increasingly used in engineering and computing education, most often through single-agent conversational interfaces that encourage rapid convergence. Such approaches offer limited support for deliberation in ill-structured problems that involve competing constraints and perspectives. This study investigates the classroom use of a multi-perspective chatbot interaction design in which students engage sequentially with multiple role-specialized AI chatbots. The activity involved three independent chatbots representing technical, institutional, and community perspectives. Each chatbot was configured with a distinct professional perspective and communication style, including limited variation in tone and assertiveness. Students interacted with the chatbots sequentially during an in-class exercise focused on two realistic infrastructure decision scenarios. Interaction order and persona tone were varied across teams to examine how these design choices influenced learner–AI dialogue. Data sources included pre-activity problem framing, chatbot interaction transcripts, in-class artifacts, post-activity questionnaires, and written reflections from twenty-four students working in twelve teams. Analysis focused on interactional patterns rather than decision outcomes or learning gains. Results indicate that engaging multiple role-specialized chatbots supported diverse deliberation trajectories, including convergence, confirmation through deliberation, and comparison across perspectives. Interaction order influenced when constraints entered reasoning, while persona tone shaped how explicitly students articulated trade-offs and justifications. Students assumed responsibility for synthesizing perspectives across chatbots, highlighting the role of learner-mediated coordination. The study contributes design-relevant insights into how role definition, interaction sequencing, and persona characteristics shape deliberative dialogue and learner–AI interaction in engineering and computing education. Findings provide practical guidance for instructors seeking to use multi-chatbot configurations to support perspective-taking and justification without complex system integration.

1. Introduction

Large language models (LLMs) are increasingly used in engineering and computing education to support problem-solving, explanation, and reflection. Most classroom deployments rely on single-agent conversational interfaces, which present learners with unified responses that often converge quickly on a single line of reasoning. While effective for tutoring and information retrieval, such systems can obscure the presence of competing constraints and values that characterize authentic engineering decision-making.

Recent work in educational AI and computer-supported collaborative learning suggests that exposure to multiple perspectives is essential for developing professional judgment in ill-structured problems. However, implementing multi-perspective dialogue in classroom settings typically requires significant instructor mediation or complex system architectures. This creates a tension between pedagogical ambition and practical deployability.

This study examines an alternative design approach using a multi-agent AI chatbot system in which learners interact sequentially with multiple role-specialized agents, each representing a distinct professional perspective. Rather than relying on adaptive orchestration or automated moderation, the system delegates coordination and synthesis to learners themselves. The instructional goal is not to optimize decisions, but to structure dialogue in ways that surface trade-offs, provoke deliberation, and require justification.

The study was conducted in a junior-level engineering economy course using two realistic infrastructure decision scenarios. Students engaged with three independently instantiated chatbot personas representing technical, institutional, and community perspectives through short, structured interactions. Interaction order and persona tone were intentionally varied to examine how design choices influence learner–AI dialogue.

The present study does not evaluate learning gains, ethical correctness, or decision quality. Engineering economy serves only as a contextual domain. The contribution lies in examining how design choices for multi-agent interaction shape deliberative reasoning processes in classroom settings.

Specifically, the study addresses the following research questions:

1. How do learners engage with and coordinate multiple AI agents representing distinct professional perspectives?
2. How does interaction sequencing influence deliberation trajectories and synthesis behaviors?
3. How does controlled variation in agent tone affect learner engagement and justification?
4. What design implications emerge for implementing multi-agent AI systems in engineering and computing education?

By answering these questions, the study contributes empirical insights into interaction design for educational multi-agent systems, offering reproducible guidance for instructors and system designers seeking to move beyond single-agent conversational tools.

Five interaction design implications emerge related to learner-mediated orchestration, role stability, sequencing as a pedagogical lever, tone modulation, and system simplicity. These implications are previewed here and elaborated in Section 7 as design-relevant contributions for educational multi-agent AI systems.

2. Background and Related Work

This study is situated at the intersection of research on online learning interactions, cognitive support for ill-structured problem-solving, and the emerging use of large language models in engineering and computing education. Rather than treating AI chatbots as instructional replacements or performance optimizers, the present work approaches interaction design as the primary instructional mechanism. Accordingly, this section synthesizes prior research that informs how interaction structure, cognitive constraints, and learner agency shape deliberative reasoning, while also identifying limitations in existing chatbot-based approaches that motivate the multi-agent design examined in this study.

Foundational theories of online learning emphasize that interaction is not merely a conduit for information delivery but a central determinant of learning quality. Anderson argues that learning environments should be understood as intentionally designed interactions among learners, instructional agents, and interfaces, with distinct forms of interaction supporting distinct cognitive processes [1]. Empirical work further demonstrates that the structure and type of interaction, rather than participation alone, significantly influence learners' perceptions of learning, engagement, and satisfaction [2]. Together, these studies establish interaction design as a legitimate object of instructional inquiry, particularly in learning contexts that emphasize discussion, judgment, and interpretation.

The role of interaction design becomes especially salient in ill-structured problems, such as ethical and economic decision-making, where learners must weigh competing constraints, values, and stakeholder perspectives. Cognitive load theory provides a useful explanatory framework for understanding why such tasks are challenging. Sweller shows that unguided problem solving can consume substantial working memory resources, leaving limited capacity for schema development and conceptual integration [3]. In these contexts, instructional structures that organize information, separate concerns, or sequence perspectives can reduce extraneous cognitive load while preserving productive challenge. From this perspective, interaction structure functions not as a simplification of the task but as a means of enabling deeper reasoning without premature convergence.

At the same time, learning sciences research cautions against instructional designs that over-scaffold or automate core aspects of learning regulation. Work on self-regulated, co-regulated, and socially shared regulation of learning emphasizes that learners must actively monitor, evaluate, and adjust their reasoning to develop durable judgment skills [4]. Extending this argument to AI-supported environments, Molenaar introduces the concept of hybrid human–AI regulation, warning that systems which assume responsibility for regulation risk undermining learners' capacity for self-regulation [5]. Hybrid approaches instead advocate designs in which AI supports reasoning processes while learners retain responsibility for coordination, reflection, and decision making.

Against this theoretical backdrop, a rapidly expanding literature examines the use of AI chatbots in engineering and computing education. Many studies frame large language models as tutors, assistants, or explainers that support content mastery, procedural problem solving, or metacognitive reflection. Reported benefits include accessibility, efficiency, and increased learner confidence, alongside concerns related to reliability, academic integrity, and overreliance [6,7]. Prior scholarship has similarly emphasized the need for revised assessment strategies, AI literacy, and intentional pedagogical integration as generative AI tools become more prevalent [8]. Earlier classroom work has also shown that custom-configured chatbots can support ethical reasoning, reflective critique, and engagement in engineering classrooms, suggesting that intentionally designed conversational roles may serve as meaningful instructional scaffolds [9]. While this literature provides important insight into adoption and engagement, most implementations rely on single-agent conversational interfaces that prioritize clarity and convergence. As a result, deliberation, perspective-taking, and judgment are often treated as downstream outcomes rather than as processes shaped by interaction design.

More recent work has begun to explore role-based and multi-agent AI systems as a means of supporting higher-order reasoning. Cao et al. demonstrate that assigning distinct roles to AI agents can structure learner inquiry [10], while Hu et al. show that persona-driven multi-agent debate frameworks can increase argumentative diversity and coherence [11]. However, these systems typically rely on automated orchestration, with the AI coordinating dialogue, sequencing perspectives, and synthesizing outputs. This design choice positions the system, rather than the learner, as the primary integrator of perspectives, implicitly framing coherence and optimization as algorithmic responsibilities.

The present work adopts a different design stance. Drawing on interaction theory, cognitive load theory, and regulation-of-learning frameworks, it treats interaction design itself as the instructional intervention.

Rather than orchestrating dialogue through a central system or synthesizing perspectives on behalf of learners, the approach assigns distinct ethical and economic perspectives to independent, role-specialized AI agents and places responsibility for coordination, comparison, and integration with learners. Role differentiation and sequential interaction structure are used to manage cognitive load, while learner-mediated orchestration preserves agency over judgment and synthesis [3-5].

Finally, this work is informed by research on student engagement, particularly distinctions between behavioral participation and deeper forms of cognitive engagement. Engagement is understood here as learners' psychological investment in reasoning, reflection, and sense-making, rather than as time-on-task or interaction frequency [12]. Recent studies of human–AI interaction patterns caution that high levels of AI use do not necessarily indicate meaningful engagement, particularly when learners adopt passive or compliance-oriented interaction strategies [13]. Accordingly, this study does not evaluate learning gains or decision quality. Instead, it examines how structured, multi-perspective AI interactions shape deliberative dialogue, justification, and learner-mediated synthesis in realistic classroom settings.

3. System and Interaction Design

This section describes the structure of the multi-agent AI chatbot system and the interaction design that supports deliberative dialogue. The emphasis is on architectural and pedagogical design choices rather than on experimental conditions or outcomes.

3.1 Multi-Agent Architecture

The system consisted of three independent chatbot instances, each configured to represent a distinct professional perspective relevant to engineering decision-making. The roles included a professional engineer, an institutional planner/manager acting as an economic steward, and a community advocate. Each chatbot operated as a standalone conversational agent, with no shared memory or inter-agent communication, ensuring that responses reflected only the assigned role, not a synthesized or negotiated position.

This architecture was designed to externalize perspective diversity rather than reconcile it algorithmically. Because agents did not synthesize their outputs, learners served as the coordination layer, interpreting, comparing, and integrating responses across agents. In this design, judgment and synthesis remained human responsibilities rather than automated functions of the system. The structural relationship among the independent agents and the learner coordination layer is illustrated in Figure 1.

Agents do not communicate or internally synthesize responses, and each interaction is bound to a single perspective at a time. Students engage agents sequentially, using insights from prior interactions to inform subsequent questioning. This structure requires learners to retain and interpret earlier responses as they construct cross-perspective comparisons and trade-off evaluations.

Coordination, comparison, and integration of perspectives are externalized to the learner, who functions as the epistemic coordination layer. By distributing perspective diversity across independent agents and externalizing integration to learners, the architecture contrasts with adaptive multi-agent systems that algorithmically reconcile viewpoints. The design treats interaction structure, not automated synthesis, as the primary instructional mechanism.

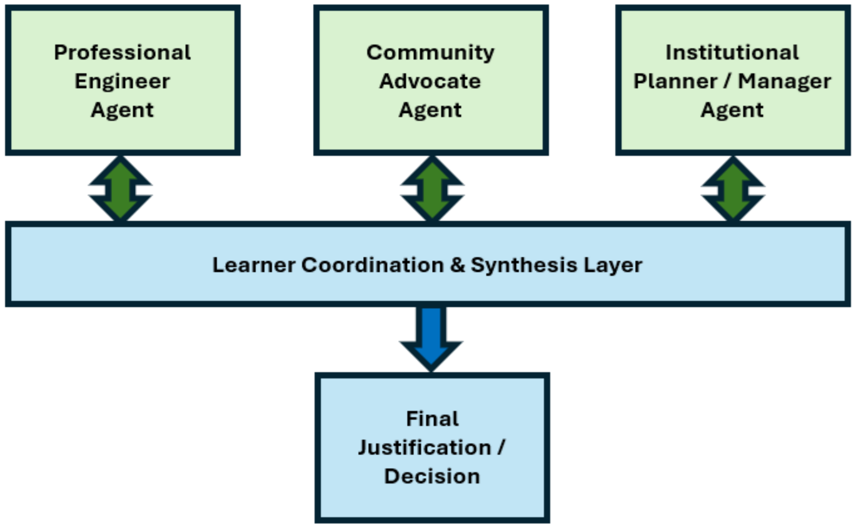

Figure 1. Learner-mediated multi-agent interaction architecture. Independent role-specialized chatbot agents provide distinct perspectives. Learners interact sequentially and perform comparison, evaluation, and synthesis. Coordination is externalized to the learner.

3.2 Persona Configuration and Role Stability

Each chatbot persona was instantiated using a structured contextual knowledge file that defined its professional background, priorities, constraints, and communication style. These configuration artifacts ensured role stability across interactions and allowed each agent to consistently foreground a distinct category of concern. Two instructional cases were used: Case A involved a campus microgrid resilience investment, and Case B focused on stormwater mitigation using permeable pavement (see Appendices A and B). Tables 1 and 2 summarize persona configurations for each case, including professional framing, primary emphasis, and interaction tone. An asterisk (*) denotes the persona assigned the strongest tone.

Table 1. Chatbot Persona Configuration and Role Emphasis for Case A (Microgrid System)

Persona Role	Professional Framing	Primary Emphasis	Interaction Tone
Professional Engineer	Licensed engineer	Feasibility, safety, technical risk	Grounded, consistent
Community Advocate	Community representative	Resilience, public impact	Engaged, value-focused
Institutional Planner	University finance & planning	Cost, trade-offs, institutional risk	*Very assertive, constraint-focused

Table 2. Chatbot Persona Configuration and Role Emphasis for Case B (Stormwater Mitigation)

Persona Role	Professional Framing	Primary Emphasis	Interaction Tone
Professional Engineer	Licensed engineer	Feasibility, sustainability	Grounded, consistent
Community Advocate	Community representative	Safety, environmental protection	*Highly passionate, urgent
Institutional Manager	Planning & capital projects	Cost, compliance, maintainability	Analytical, cautious

Role priorities, including technical feasibility for the engineer, fiscal responsibility for the institutional planner/manager, and stakeholder impact for the community advocate, remained fixed throughout the activity. Limited variation in tone and assertiveness was introduced at the case level. This separation between role commitments and rhetorical framing enabled examination of tone effects without conflating perspective content with stylistic variation. Stable role boundaries supported interpretability during interaction, allowing students to attribute arguments to identifiable epistemic positions and engage in structured comparison and trade-off reasoning across agents.

3.3 Interaction Structure and Learner Orchestration

Students engaged with the chatbot personas sequentially in structured, time-bounded exchanges. Each team interacted with one agent at a time, posing questions, reviewing responses, and documenting emerging insights before proceeding to the next perspective. The approach did not include automated moderation, synthesis, or cross-agent coordination. Sequential interaction was selected to foreground one perspective at a time and to reduce cognitive overload in ill-structured decision contexts. By isolating perspectives during engagement, the design encouraged focused questioning and deliberate comparison rather than immediate convergence toward a unified response. This structure required students to retain and interpret prior agent inputs as they prepared questions for subsequent interactions.

Because agents did not communicate with one another, integration occurred exclusively through learner reflection and synthesis. Students were responsible for identifying tensions, articulating trade-offs, and developing a final justification based on multiple perspectives. The interaction structure, therefore, positioned dialogue as an epistemic process that unfolded through structured exposure to distinct professional viewpoints. This approach positions interaction design as the primary instructional mechanism. Rather than directing students toward a predefined resolution, the system was intended to surface competing constraints and require explicit reasoning about feasibility, responsibility, and stakeholder impact.

3.4 Design Rationale

The system was designed to prioritize simplicity, transparency, and instructional robustness. Rather than relying on specialized orchestration software or inter-agent communication protocols, the architecture used independent chatbot instances configured through structured persona definitions. This approach ensured that system behavior remained visible and interpretable to both instructors and students. A central design decision was to externalize coordination to learners rather than automate synthesis within the system. Many multi-agent AI frameworks emphasize algorithmic aggregation or negotiated resolution among agents. In contrast, this design treated diversity of perspective as an instructional resource and positioned learners as the locus of comparison, evaluation, and integration. Separating perspective generation from reconciliation preserved human judgment as the decisive element in the deliberative process. The design also reflects pedagogical commitments drawn from deliberative learning theory and learning regulation frameworks. Structured exposure to competing constraints was intended to prompt articulation of trade-offs rather than rapid convergence. Stable role identity, controlled variation in tone, and sequential interaction were used to support interpretability and sustained reasoning in ill-structured problems. Finally, the architecture was developed for practical classroom deployment. The system can be implemented using standard large language model interfaces without additional infrastructure, enabling replication across engineering and computing education contexts. Emphasizing structural clarity over algorithmic complexity reinforces the instructional rationale for externalizing coordination to learners.

4. Study Context and Instructional Deployment

This section describes the instructional setting, decision cases, and classroom integration of the multi-agent system. The purpose is to establish context for interpreting interactional findings rather than to evaluate domain-specific learning.

4.1 Course Context

The study was conducted in a junior-level engineering economy course. The activity was integrated into regular course instruction and completed during a single class session. Students worked in pairs, reflecting collaborative practices used throughout the course. The exercise was designed as an applied decision activity rather than a graded assessment. Students were evaluated on completion and reflection rather than on the correctness of their decisions. This structure emphasized deliberative engagement and justification over the accuracy of outcomes.

4.2 Decision Cases

Two realistic infrastructure decision cases were used to ground the activity. The cases followed the National Society of Professional Engineers' (NSPE) ethics case study format, which presents a contextual vignette featuring an ethical dilemma intersecting with an engineering decision. In this investigation, the format was adapted to problems situated on campus and in the surrounding community. Case A involved a campus microgrid investment proposal to enhance resilience during power disruptions. Case B focused on stormwater management and the use of permeable pavement near a flood-prone watershed.

Summary descriptions of the cases are provided in Appendices A and B. While the technical details differed, both cases were structured to surface competing constraints related to feasibility, institutional resources, and community impact. Each case presented plausible alternatives with associated trade-offs, creating conditions for engagement with multiple professional viewpoints.

4.3 Classroom Deployment

Students began the activity by documenting initial problem framing and tentative decision rationale. Teams then engaged sequentially with the three chatbot personas. Following completion of all interactions, students submitted a written justification outlining their final recommendation and the reasoning used to reconcile competing perspectives. An individual questionnaire and short reflection were completed after the team exercise. The deployment emphasized structured exposure to distinct professional viewpoints while preserving student responsibility for coordination and synthesis.

4.4 Scope and Intent

The instructional deployment was designed to explore how learners interact with and coordinate multiple AI agents. The study does not evaluate instructional effectiveness or compare decision quality across treatments. Engineering economy and ethics provided contextual grounding, but the analytic focus remained on interaction design, dialogue patterns, and learner-mediated synthesis.

5. Methods and Analysis

This section describes the study design, participant structure, interaction conditions, and analytical approach used to examine learner–agent dialogue patterns.

5.1 Participants and Study Design

A total of 12 student teams participated in the activity, with six teams assigned to each case. Students worked in pairs, resulting in 24 individual participants. The activity was conducted during a single class session of a regularly scheduled course. The study employed a structured interaction design in which teams engaged sequentially with three independent chatbot personas. Interaction sequencing and tone emphasis were varied across teams to examine how the timing and framing of perspective exposure influenced deliberative processes. The study was not designed to evaluate decision correctness or learning gains, but rather to examine interaction patterns and synthesis behaviors.

5.2 Interaction Conditions

In each case, teams were assigned to one of two interaction-sequence conditions. Sequencing determined the order in which students encountered the engineer, institutional planner/manager, and community advocate personas. Tone variation was introduced at the case level to foreground dominant contextual tensions while preserving stable role commitments as described in Section 3.

Table 3 summarizes the interaction conditions across cases, including sequencing structure and tone configuration. The "treatment" applied is denoted with an asterisk.

Table 3. Interaction Sequencing and Tone Conditions Across Cases

Case	Teams (n)	Sequence	Agent Interaction Order
A	3	A1	Engineer → Community Advocate → *Planner
A	3	A2	*Planner → Community Advocate → Engineer
B	3	B1	Engineer → *Community Advocate → Manager
B	3	B2	Manager → *Community Advocate → Engineer

Following interaction with the agents, teams produced a written justification outlining their final recommendation and the reasoning process used to reconcile competing perspectives.

The study did not evaluate the correctness of decisions or learning gains. Instead, the analysis focused on interaction patterns, synthesis behaviors, sequencing effects, and tone-related dynamics observed during deliberation.

5.3 Activity Structure and Data Sources

Multiple qualitative data sources were collected to support analysis of learner–agent interaction patterns. The intent was to capture both process-level dialogue and outcome-level synthesis artifacts. Table 4 summarizes the data sources collected for each team, including interaction transcripts, written artifacts, and post-activity reflections. Full transcripts of learner–agent interactions provided process-level evidence of questioning patterns, sequencing effects, and tone-related dynamics. Team worksheets captured initial framing and evolving problem definitions. Final written justifications documented synthesis behaviors and articulated trade-offs. Individual questionnaires and reflections offered insight into perceived engagement and reasoning strategies. Collectively, these artifacts enabled cross-case comparison of interaction trajectories and supported triangulation of observed patterns.

Table 4. Data Sources Collected for Analysis

Data Source	Collected From	Purpose in Analysis
Pre-activity worksheet	Individuals	Initial framing and baseline reasoning
Chatbot interaction transcripts	Teams	Dialogue progression and questioning behavior
In-class worksheet notes	Teams	Real-time synthesis and perspective uptake
Post-activity questionnaire	Individuals	Perceived influence and engagement
Written reflections	Individuals	Final justification and synthesis behavior

5.4 Analytical Approach

Analysis focused on identifying recurring interaction patterns, sequencing effects, synthesis behaviors, and tone-related dynamics within learner–agent dialogue. The primary unit of analysis was the team-level interaction sequence rather than isolated individual utterances. Coding proceeded in two stages and involved two reviewers, including the instructor-researcher. During the open-coding phase, reviewers independently identified emergent interaction behaviors across transcripts and written artifacts, then met to compare interpretations and reconcile category definitions. This stage emphasized descriptive categorization of how teams framed problems, engaged agent perspectives, and articulated trade-offs. In the second stage, focused coding examined relationships among sequencing conditions, tone variation, and synthesis behaviors to identify patterned differences in deliberation trajectories. Artifacts from Table 4 were analyzed in combination to support triangulation. Interaction transcripts provided process-level evidence of questioning and response dynamics, while final written justifications and reflections provided evidence of synthesis and perspective integration. Patterns were compared across cases to identify convergent and divergent deliberative structures. The study was not designed for causal inference. Findings are presented as descriptive interaction patterns that inform instructional design considerations for multi-agent educational systems.

5.5 Validity and Limitations

Several limitations constrain the interpretation of the findings. First, the study was conducted within a single course and institution, with a limited number of teams. The sample size supports exploratory pattern identification but does not permit statistical generalization. Second, it was not designed to evaluate learning gains or instructional effectiveness. The analytic focus was on interaction structure and deliberative behavior rather than on outcome improvement. As such, the findings inform design implications rather than claims about comparative instructional impact. Third, tone and sequencing were manipulated at a coarse level and within bounded contextual cases. Additional research is required to examine how these design variables function across different disciplines, problem types, or levels of student experience. Finally, because learners served as the coordination layer, variations in prior experience, confidence, and team dynamics may have influenced deliberative trajectories. These contextual influences were not experimentally controlled. Despite these limitations, the structured deployment and multi-source data collection support credible identification of recurring interaction patterns relevant to educational multi-agent design.

6. Results

This section reports observed interactional patterns across teams engaging with the multi-agent AI chatbot system. Findings are qualitative and exploratory, emphasizing how interaction design features such as role differentiation, sequencing, and agent tone shaped learner–AI dialogue and deliberation. Results are organized around deliberation trajectories, sequencing effects, emergent synthesis behaviors, and tone-related interaction dynamics, rather than decision outcomes or learning gains.

Table 5 provides an overview of the empirical basis for the findings reported in this section. The patterns summarized in the table were identified through analysis of team worksheets, chatbot interaction transcripts, in-class artifacts, and post-activity reflections, as described in Section 5, Methods and Analysis. The table highlights how frequently specific interaction behaviors occurred across cases and sequencing conditions. Counts are presented at a descriptive level to support interpretation rather than statistical generalization. The subsections that follow elaborate on these patterns with contextual detail and illustrative examples.

Table 5. Summary of Observed Interactional Patterns Across Teams

Observed Interactional Pattern	Evidence Across Teams
Multi-perspective references in the final justification	Observed in 9 of 12 teams
Justification refinement without decision change	Observed in 7 of 12 teams
Early constraint anchoring (planner/manager-first)	Observed in 6 of 6 teams
Conditional or phased justification	Observed in 6 of 12 teams
Explicit cross-perspective comparison	Observed in 8 of 12 teams
Explicit challenge of agent responses (Case A)	Observed in 5 of 6 teams
Explicit challenge of agent responses (Case B)	Observed in 4 of 6 teams

6.1 Deliberation Trajectories Across Teams

In both cases, student teams exhibited distinct trajectories of deliberation as they sequentially engaged with multiple chatbot personas. These trajectories describe how reasoning evolved during interaction rather than whether teams altered their final decisions.

One trajectory involved convergence through deliberation. Teams refined initial problem framing as new perspectives were introduced, reconsidering assumptions, expanding evaluative criteria, or modifying preferred options. These teams frequently referenced multiple agent perspectives explicitly in their final justifications. A second trajectory involved confirmation through deliberation. Teams maintained an initial preference but used interaction to validate and strengthen their reasoning. Rather than simply reaffirming their position, they tested it against alternative perspectives and articulated why certain constraints were deprioritized. A third trajectory involved comparison through deliberation. Teams treated agent interactions as parallel evaluative inputs, emphasizing contrasts among perspectives without fully reconciling them. These teams highlighted trade-offs while maintaining less definitive commitment to a dominant rationale. These trajectories were observed across both cases and sequencing conditions, suggesting that deliberative patterns emerged from interaction structure rather than case content alone.

6.2 Effects of Interaction Sequencing

Interaction sequencing influenced how deliberation unfolded, particularly when constraints entered. Teams encountering constraint-oriented perspectives early more frequently incorporated feasibility, affordability, and institutional considerations into subsequent questioning. Early exposure anchored later dialogue, prompting targeted questions about mitigation strategies and trade-offs. Teams encountering constraints later began with broader exploration before narrowing focus. In these cases, constraints functioned as a refining influence rather than an initial anchor, shaping synthesis at later stages. Sequencing effects were also evident in interactions with the community advocate. When preceded by a constraint-focused agent, students framed community concerns in terms of feasibility and institutional capacity. When preceded by the engineer, interactions more often emphasized normative priorities and societal impact before addressing constraints. These results suggest interaction order influenced the structure and timing of deliberation, even when final decisions remained stable. Sequencing shaped how students framed questions and integrated perspectives, reinforcing its role as a meaningful design parameter in multi-agent educational systems.

6.3 Emergent Synthesis Behaviors Across Agents

As students engaged with multiple agents, several synthesis behaviors emerged. These patterns arose as students assumed responsibility for integrating perspectives. One pattern, perspective stacking, involved retaining multiple agent viewpoints as distinct evaluative layers without immediate resolution. Students attributed concerns to identifiable roles and deferred reconciliation, enabling sustained deliberation. A second pattern involved arbitration. Students weighed competing perspectives to determine priority, often referencing feasibility, urgency, or institutional constraints when resolving tensions. A third pattern involved reconciliation. Students attempted to integrate opposing perspectives into conditional or hybrid rationales, such as phased implementation strategies or contingent recommendations that acknowledged unresolved trade-offs.

Synthesis behaviors were influenced in part by sequencing. Teams exposed early to strongly framed constraints were more likely to engage in arbitration, whereas those exposed later to constraints were more likely to pursue reconciliation. However, synthesis remained learner-driven, and no single pattern dominated. Together, these patterns indicate that multi-agent interaction supports active coordination and perspective integration without automated moderation. Synthesis functioned as a process of structured justification rather than evidence of decision change.

6.4 Influence of Agent Tone and Assertiveness

In addition to sequencing, tone and assertiveness influenced how students engaged with agent perspectives. Tone variation was introduced as a design parameter aligned with dominant tensions, while substantive role priorities remained stable. This variation shaped how strongly particular constraints or values were foregrounded during interaction.

In Case A, the institutional planner adopted an assertive tone emphasizing fiscal pressure, deferred maintenance, and trade-offs associated with capital investment. In Case B, the community advocate used a more urgent and passionate tone, emphasizing public safety, environmental protection, and human consequences. The engineer persona maintained a consistent, technically grounded tone across conditions.

The following excerpt (from Case A) illustrates how assertive framing prompted evaluative engagement:

- Student: “If we pursue the solar-plus-battery option, how do we justify the additional cost given deferred maintenance pressures?”
- Institutional Planner (assertive tone): “Allocating substantial capital to battery storage under current fiscal constraints risks delaying critical infrastructure upgrades and may amplify institutional financial exposure.”
- Student (follow-up): “Would a phased implementation or partial storage capacity reduce that exposure while still addressing outage risk?”

In this exchange, assertive framing foregrounded fiscal risk and opportunity cost. Rather than deferring to the recommendation, the student reframed the constraint as a design variable and proposed a conditional modification. Similar patterns across teams suggest that an assertive tone increased the salience of constraints and prompted evaluative questioning without producing uncritical agreement. Teams engaging with strongly framed agents more frequently articulated conditional reasoning and explicitly acknowledged trade-offs. In contrast, interactions with more measured agents produced technically detailed but less explicitly value-oriented justifications. Tone also interacted with sequencing. Early exposure to an assertive agent tended to anchor subsequent deliberation, whereas later exposure more often prompted reconciliation or retrospective justification. Across teams, students regularly questioned or contextualized strongly framed arguments, indicating that assertiveness increased engagement without diminishing learner agency.

7. Discussion and Design Implications

This section interprets observed interaction patterns as design implications for educational multi-agent AI systems in engineering and computing education. The discussion emphasizes how roles, sequencing, tone, and learner orchestration shape deliberative dialogue.

7.1 Design Implications for Educational Multi-Agent AI Systems

Drawing on the interaction patterns reported in Sections 6.1 through 6.4, five design implications emerge for implementing multi-agent AI systems in engineering and computing education. Each implication is grounded in observed learner–AI dialogue behaviors and intended to support replication across instructional contexts.

<u>Implication 1: Learners Can Function as the Coordination Layer</u> - The findings suggest that learner-mediated orchestration can support coordination of multi-agent interaction when roles are clearly differentiated, and interaction windows are structured. Students integrated perspectives through synthesis behaviors such as stacking, arbitration, and reconciliation (Section 6.3). Externalizing coordination to learners preserved agency and intensified deliberative engagement without automated moderation.

<u>Implication 2: Role Stability Is More Critical Than Adaptive Intelligence</u> - Students treated each chatbot as a stable epistemic source for a distinct category of concern, enabling structured comparison and trade-off reasoning. In contrast to adaptive systems that dynamically synthesize perspectives, this design prioritized consistent role boundaries. Stability, rather than conversational flexibility, emerged as central to supporting deliberation.

<u>Implication 3: Sequencing Functions as a Pedagogical Lever</u> - The order in which students encountered agent perspectives influenced how deliberation unfolded (Section 6.2). Early exposure to constraint-

oriented agents anchored subsequent questioning, whereas later exposure prompted reconciliation and conditional justification. Sequencing functions as a temporal mechanism shaping deliberation without altering case content.

Implication 4: Persona Tone Modulates Deliberative Intensity - Variation in tone influenced the salience of concerns (Section 6.4). Assertive framing increased evaluative questioning and trade-off articulation without diminishing learner agency. Tone serves as a design lever that can amplify deliberative intensity while preserving human judgment.

Implication 5: Simplicity and Transparency Support Adoption and Responsible Use - The system relied on independent chatbot instances and explicit role definitions rather than complex orchestration. This simplicity enhanced classroom feasibility and interpretability. Clear role boundaries and visible learner coordination reduced the risk of overreliance and supported responsible AI deployment.

These implications suggest that effective educational multi-agent systems need not rely on complexity. Structured roles, sequencing, tone variation, and learner-mediated synthesis can support deliberative reasoning.

7.2 Comparison of Single-Agent Educational Chatbots

This study did not include an experimental comparison between multi-chatbot and single-agent configurations. Accordingly, the contrasts discussed here are design-informed rather than performance-based. Observed interactional patterns provide insight into how multi-chatbot configurations differ from more common single-agent use. Single-agent chatbots typically present a unified conversational voice that synthesizes perspectives internally. In instructional contexts, this approach supports efficiency and clarity, but it may also encourage rapid convergence on a single line of reasoning. These findings suggest that educational AI system design may benefit from architectures that make competing perspectives visible to learners rather than resolving them internally before interaction begins. By presenting an integrated response, the system reduces the need for learners to reconcile competing viewpoints explicitly. Prior work using a single-agent chatbot suggested that customized AI advisors can support student engagement and ethical reasoning [9]. The present study extends that line of inquiry by externalizing perspective diversity across multiple coordinated agents rather than embedding guidance within a single conversational voice.

In contrast, the multi-chatbot interaction design examined in this study externalized perspective diversity by distributing distinct concerns across role-specialized agents. Learners engaged each perspective independently and reconciled tensions through deliberation. Observed interactional patterns suggest that students treated chatbot responses as partial and situated inputs rather than as authoritative answers. Teams frequently referenced specific roles when articulating justifications, explicitly compared perspectives, and acknowledged unresolved trade-offs.

These behaviors align with deliberative reasoning processes in which perspective integration and trade-off articulation are central. However, this analysis does not imply superiority of one configuration over another. Single-agent chatbots may be well-suited for instructional tasks emphasizing explanation, practice, or procedural guidance. The findings presented here suggest that multi-chatbot configurations offer a complementary approach to interaction design for contexts in which perspective-taking, justification, and the negotiation of competing constraints are instructional priorities. Unlike AI-moderated debate systems that algorithmically reconcile competing viewpoints, the present design deliberately externalizes epistemic tension to the human coordination layer. Learners, rather than the system, function as the integrators of meaning.

7.3 Responsible AI Considerations and Limitations

The interpretation of these findings is shaped by several constraints. The study involved a limited number of teams within a single course and institution, and the interaction window was brief. Personas were instructor-designed and non-adaptive, and the activity was not embedded in a longitudinal sequence. These conditions support exploratory pattern identification rather than generalization. From a design perspective, care is required in constructing persona roles to avoid embedding unintended bias or reinforcing institutional hierarchies through tone or framing. Although assertiveness increased evaluative engagement in this study, excessive authority signaling in other contexts could discourage dissent. Accordingly, transparency, role clarity, and non-authoritative positioning remain essential. Externalizing coordination to learners helps preserve human judgment and reduces reliance on automated synthesis in educational AI systems.

7.4 Extending Interaction Design Toward Learning and Assessment Contexts

The present study isolates interaction design as the primary instructional intervention and does not evaluate learning gains or decision quality. However, the structured multi-agent configuration provides a foundation for future work that connects interaction design with the assessment of reasoning processes. Because perspectives are distributed across role-stable agents and coordination is externalized to learners, the system makes trade-off articulation and justification visible. Future research may examine argument quality, perspective integration, and longitudinal changes in reasoning through repeated exposure. Experimental variation in sequencing and tone could further clarify their influence on deliberative development. This work establishes a foundation for connecting structured multi-agent dialogue with assessment in engineering and computing education.

7.5 Conclusion

This study examined how a learner-mediated, multi-agent AI chatbot configuration can shape deliberative dialogue in engineering education contexts. Rather than functioning as tutors or decision authorities, the chatbots operated as role-stable advisors that surfaced distinct professional perspectives. Coordination and synthesis were intentionally externalized to learners, positioning students as the integrators of meaning. Observed interaction patterns indicate that role differentiation, sequencing, and calibrated variation in tone influenced how students framed questions, articulated trade-offs, and justified decisions. Deliberation unfolded through convergence, confirmation, and comparison trajectories, with synthesis behaviors emerging as students reconciled competing constraints. Tone modulation increased deliberative intensity without diminishing learner agency, and sequencing shaped when and how constraints entered reasoning.

Collectively, the findings suggest that carefully structured multi-agent interaction can serve as a dialogic infrastructure for perspective-taking and trade-off analysis. By prioritizing role stability, transparent orchestration, and intentional design variation, educational multi-agent systems can support structured reasoning while preserving human judgment. These implications extend beyond engineering economy to broader engineering and computing education contexts, where students must navigate competing constraints and construct defensible justifications.

References

[1] T. Anderson, *Towards a theory of online learning. Theory and practice of online learning*, vol. 2. 2004.

[2] J. B. Arbaugh and R. Benbunan-Fich, “The importance of participant interaction in online environments,” *Decis. Support Syst.*, vol. 43, no. 3, pp. 853–865, Apr. 2007, doi: 10.1016/j.dss.2006.12.013.

[3] J. Sweller, “Cognitive Load During Problem Solving: Effects on Learning,” *Cogn. Sci.*, vol. 12, no. 2, pp. 257–285, Apr. 1988, doi: 10.1207/s15516709cog1202_4.

[4] A. F. Hadwin, S. Järvelä, and M. Miller, *Self-regulated, co-regulated, and socially shared regulation of learning. In B. J. Zimmerman & D. H. Schunk (Eds.), Handbook of self-regulation of learning and performance*. Routledge/Taylor & Francis Group. [Online]. Available: https://www.researchgate.net/publication/385592933_Self-regulated_co-regulated_and_socially_shared_regulation_of_learning_in_collaborative_learning_environments

[5] I. Molenaar, “The concept of hybrid human-AI regulation: Exemplifying how to support young learners’ self-regulated learning,” *Comput. Educ. Artif. Intell.*, vol. 3, p. 100070, 2022, doi: 10.1016/j.caeai.2022.100070.

[6] L. Benotti, M. C. Martínez, and F. Schapachnik, “Engaging high school students using chatbots,” in *Proceedings of the 2014 conference on Innovation & technology in computer science education - ITiCSE ’14*, Uppsala, Sweden: ACM Press, 2014, pp. 63–68. doi: 10.1145/2591708.2591728.

[7] Fauziah, Novita Diana, Silvia Putri, and Teuku Fadhli, “Student Engagement in English Language Learning with Artificial Intelligence (AI) Chatbot from Theory of Planned Behavior Perspective,” *JEELS J. Engl. Educ. Linguist. Stud.*, vol. 12, no. 2, pp. 585–614, Jun. 2025, doi: 10.30762/jeels.v12i2.4085.

[8] R. Smith and H. Lester, “Navigating the AI Revolution in Engineering Management Education: Strategies for Detection, Integrity, and Pedagogical Enhancement,” in *2024 ASEE Annual Conference & Exposition Proceedings*, Portland, Oregon: ASEE Conferences, Jun. 2024, p. 47796. doi: 10.18260/1-2--47796.

[9] R. Smith, R. Castles, and E. Sondergard, “The Engineering Ethics Coach: Integrating AI Chatbots to Strengthen Ethical Decision-Making in Engineering Economy,” in *2025 ASEE Annual Conference & Exposition Proceedings*, Montreal, Quebec, Canada: ASEE Conferences, Jun. 2025, p. 57218. doi: 10.18260/1-2--57218.

[10] C. C. Cao, Z. Ding, J. Lin, and F. Hopfgartner, “AI Chatbots as Multi-Role Pedagogical Agents: Transforming Engagement in CS Education,” Aug. 08, 2023, *arXiv*: arXiv:2308.03992. doi: 10.48550/arXiv.2308.03992.

[11] Z. Hu, H. P. Chan, J. Li, and Y. Yin, “Debate-to-Write: A Persona-Driven Multi-Agent Framework for Diverse Argument Generation,” 2024, *arXiv*. doi: 10.48550/ARXIV.2406.19643.

[12] J. Fredricks, P. Blumenfeld, and A. H. Paris, “School Engagement: Potential of the Concept, State of the Evidence,” *Rev. Educ. Res.*, vol. 74, no. 1, pp. 59–109, Mar. 2004, doi: 10.3102/00346543074001059.

[13] X. Wang *et al.*, “What matters in AI-supported learning: A study of human-AI interactions in language learning using cluster analysis and epistemic network analysis,” *Comput. Educ.*, vol. 194, p. 104703, Mar. 2023, doi: 10.1016/j.compedu.2022.104703.

APPENDIX A

Case A: Campus Microgrid for Resilience and Sustainability Context

Case Overview

In Case A, student teams evaluated two alternative designs for a proposed campus microgrid to improve power resilience at East Carolina University (ECU) in Greenville, North Carolina. The decision context focused on maintaining electrical service to critical facilities, including the ECU School of Medicine, student housing, and research laboratories, during extended regional power outages caused by hurricanes and severe weather events. Multi-Perspective

At the time of the study, ECU relied entirely on Greenville Utilities for electrical service. While sufficient for normal operations, the regional grid has historically failed to maintain reliable power during major storms, posing risks to public safety, research integrity, and institutional operations. The proposed microgrid was framed as a long-term infrastructure investment that could reduce dependency on the regional grid while advancing sustainability goals.

Regional and Institutional Context

Eastern North Carolina is highly vulnerable to hurricanes and tropical storms, and ECU has experienced repeated power outages and flooding. Hurricane Floyd (1999) caused catastrophic flooding in the Tar River basin, resulting in an 11-day campus closure and prolonged power outages. More recent storms, including Hurricanes Matthew (2016) and Florence (2018), produced additional outages and operational disruptions, reinforcing the university's exposure to regional grid failures. These events have shaped ECU's infrastructure planning priorities. Power continuity for medical, housing, and research facilities is viewed as critical to the university's mission. At the same time, ECU faces a substantial deferred-maintenance backlog of more than $80 million, creating ongoing tension between resilience investments and other capital needs.

Design Options Under Consideration

Student teams evaluated two alternative system designs:

1. Solar-Only System - This option consisted of an 8-MW solar photovoltaic array connected to the existing utility grid. While it reduced long-term electricity purchases and supported sustainability objectives, the system could not operate independently during a grid outage. As a result, it provided no backup power for critical facilities when regional service failed. The system had a lower initial cost and lower annual operating and maintenance expenses.
2. Solar Plus Battery Storage System - This option combines the same 8-MW solar array with 20 MWh of battery storage, enabling islanded operation during grid outages. The system could supply power to designated critical loads during extended outages, significantly improving campus resilience. However, it required substantially higher initial investment and higher annual operating and maintenance costs, potentially delaying other capital projects or increasing financial pressure on the institution.

Economic and Risk Considerations

The expected project life for both options was 25 years, with financial planning conducted using an institutional interest rate of 8 percent. Annual electricity savings were estimated at approximately $1.1 million for the solar-only system and $1.25 million for the solar-plus-battery system, reflecting modest additional value from peak shaving and limited backup capability.

A key differentiator between the options was the reduction of risk. Major hurricane events causing extended outages occur roughly once every 5 to 7 years, with estimated avoided losses of approximately $2.5 million per event for critical housing, medical, and research operations. These avoided losses were uncertain but potentially significant, especially given the historical impacts of prolonged outages on campus operations.

Stakeholder Perspectives Represented by AI Chatbots

During the in-class activity, student teams conducted short interviews with three AI-based subject-matter experts, each representing a distinct institutional or community perspective:

- Jordan Taylor, P.E., Senior Engineer in ECU Facilities, emphasized professional responsibility, system reliability, public safety, and compliance with engineering ethics.
- Maria Lopez, Community Advocate and liaison to ECU, focused on sustainability, environmental stewardship, and equity considerations affecting the broader community.
- Chris Reynolds, University Associate Vice Chancellor for Facilities Planning, emphasized budget discipline, opportunity cost, deferred maintenance pressures, and institutional optics.

Teams engaged with these experts in varying sequences and were required to reconcile competing technical, economic, and ethical considerations to develop a defensible recommendation.

Disclaimer: This case is a fictionalized instructional scenario created for educational purposes. Named organizations and regional events provide contextual grounding, but the project circumstances, decision alternatives, characters, roles, and stakeholder positions are fictional and do not represent any actual person, position, institutional decision, project, or event.

APPENDIX B

Case B: Stormwater Runoff and Permeable Pavement Decision Context

Case Overview

In Case B, student teams evaluated stormwater management options for a proposed parking expansion at East Carolina University (ECU), located in Greenville, North Carolina, near the athletic complex and adjacent student housing. The site is located within the Tar River watershed, an area with a documented history of flooding during major storm events. The decision context required teams to balance cost, flood mitigation, environmental sustainability, and professional responsibility in recommending an appropriate design approach.

The expansion of impervious surfaces associated with new parking infrastructure raised concerns about increased stormwater runoff, pollutant transport, and downstream flooding impacts. These concerns were heightened by the university's proximity to flood-prone areas and by regional experiences with severe weather events.

Regional and Institutional Context

Eastern North Carolina has experienced repeated flooding events that disrupted campus operations and surrounding communities, including major impacts during Hurricanes Matthew (2016) and Florence (2018). These events highlighted the vulnerability of existing infrastructure and the potential consequences of inadequate stormwater control.

At the institutional level, ECU faces competing capital demands, including an extensive deferred-maintenance backlog and aging utility systems. Funding for the parking project would draw from the same capital pool used for other critical infrastructure upgrades, creating opportunity costs associated with higher initial investment.

At the regulatory level, the North Carolina Department of Environmental Quality is reviewing potential updates to stormwater management standards that could become mandatory within several years, raising the possibility of future retrofit requirements for systems that only meet current minimum compliance.

Design Options Under Consideration

Teams evaluated two design alternatives:

1. Minimum Compliance Option - This option relied on traditional asphalt paving and standard retention ponds, designed to meet existing stormwater regulations. It had the lowest initial cost and used established construction and maintenance practices. However, it provided limited runoff reduction and pollutant removal and could exacerbate downstream flooding risks under extreme weather conditions.
2. Enhanced Sustainability Option - This option incorporated permeable pavement, bioswales, and vegetated drainage features. The system exceeded current regulatory requirements and was intended to reduce runoff volume, improve water quality, and enhance flood resilience. While offering environmental benefits, this option entailed higher initial and ongoing maintenance costs and introduced uncertainty about long-term performance, particularly given documented maintenance challenges at similar regional installations.

Economic and Risk Considerations

The expected project life was 25 years, with planning conducted using an institutional interest rate of 8 percent. The enhanced sustainability option was associated with estimated flood-damage avoidance benefits ranging from \$50,000 to \$90,000 per year, with major flood events occurring approximately once every 6 to 8 years and causing damages on the order of \$500,000 to \$1 million per event.

Maintenance reliability emerged as a key source of uncertainty. Several municipalities in eastern North Carolina have reported performance degradation in permeable pavement systems due to sediment clogging and limited staffing for upkeep, raising questions about whether projected benefits could be sustained over time.

Stakeholder Perspectives Represented by AI Chatbots

During the in-class activity, student teams conducted short interviews with three AI-based subject-matter experts, each representing a distinct perspective:

- Dana Whitaker, P.E., Professional Engineer for ECU Facilities, emphasized technical defensibility, public safety, regulatory compliance, and professional responsibility.
- Tanya Greene, Community Advocate, focused on flood impacts, ecosystem health, and moral obligations to protect surrounding neighborhoods and waterways.
- Eric Dalton, ECU Project Manager, highlighted budget constraints, opportunity costs, regulatory risk, and alignment with institutional priorities.

Teams engaged with these experts in varying sequences and were required to reconcile conflicting viewpoints to develop a defensible recommendation.

Disclaimer: This case is a fictionalized instructional scenario created for educational purposes. Named organizations and regional events provide contextual grounding, but the project circumstances, decision alternatives, characters, roles, and stakeholder positions are fictional and do not represent any actual person, position, institutional decision, project, or event.