ProDAIS: GenAI-Facilitated Collaborative Learning System for Productive Group Discussion in Information Technology Education

Zihan Zhang[0009-0006-1817-6284], Flora Ji-Yoon Jin[0009-0002-1825-4028], Delvin Varghese[0000-0002-7806-1001], and Jiazhou Liu[0000-0003-4541-0233]

Monash University, Victoria, Australia

Abstract. Group discussion plays an essential role in learning processes, enabling students to exchange ideas, negotiate meaning, and co-construct knowledge through interaction. However, sustaining such productive interaction remains a persistent challenge. Students often experience regulatory breakdowns in how they manage participation, focus, and reflection. For example, uneven participation, topic drift, and limited joint review of progress, which hinder shared reasoning and collective outcomes. Recent advances in Generative AI (GenAI) offer new opportunities to augment these processes, not by replacing human facilitator, but by acting as a low-interference regulatory partner to support group discussion in real time. To explore this potential, the present study adopted a Design-Based Research (DBR) approach consisting of two phases. Phase 1 involved semi-structured interviews with 15 Information Technology (IT) students to examine their collaborative experiences and expectations of AI support. The findings revealed that students prioritised support for the often invisible conditions of collaboration (e.g., pacing and participation equity) rather than content generation. Guided by these needs, we developed Productive Discussion AI System (ProDAIS), a lightweight GenAI facilitator designed to support the Four Goals of Productive Discussions (FGPDs). ProDAIS integrates Socially Shared Regulation of Learning (SSRL) to determine when and why to intervene (e.g., regulating participation balance), and Academically Productive Talk (APT) to determine what and how to intervene through non-evaluative, inviting talk moves. In Phase 2, the study focused on the iterative design and evaluation of ProDAIS through a user study, examining perceived usability (System Usability Scale, SUS), cognitive load (NASA-TLX), and students' perceived support toward the FGPDs goals. Results indicate high perceived usability (SUS score: 78.5), low cognitive load, and consistent perceived support for participation balance and time management. Overall, the findings suggest that ProDAIS demonstrates how GenAI can facilitate productive group discussion by making participation balance and pacing visible.

Keywords: Generative AI · GenAI facilitation · Collaborative learning · IT education · Computer-Supported Collaborative Learning · Productive group discussion · Socially Shared Regulation of Learning.

1 Introduction

Group discussion is a central mechanism of collaborative learning, enabling learners to construct knowledge, communicate effectively, and develop teamwork skills through shared reasoning [1, 2]. In IT education, collaborative discussion is important, as students frequently work together on programming, design, and problem-solving tasks that require joint reasoning and coordination. However, sustaining productive discussion in collaborative settings remains challenging. Groups often experience uneven participation, prolonged silence, topic drift, or the absence of shared conclusions [3, 4]. As a result, discussions may remain at a surface level—focused on exchanging information or dividing tasks—rather than supporting shared sense-making and knowledge co-construction, which are essential for fostering critical thinking and deeper cognitive engagement [5, 6].

Advances in educational technologies have sought to address these challenges through computer-supported collaborative learning (CSCL), learning analytics, and AI-embedded systems [7, 8]. More recently, progress in GenAI has enabled context-aware support through natural language understanding and generation, creating new opportunities for facilitating group dialogue. However, many existing AI-facilitated discussion systems continue to prioritise task coordination, content generation, or basic technical support, with relatively limited attention to interaction quality and regulation [9, 10]. Without strong theoretical grounding or learner-informed design, such systems risk promoting superficial participation or passive dependence, rather than fostering deeper reasoning and peer engagement. Moreover, many AI-facilitated discussion tools are designed *for* students rather than *with* them, relying on fixed facilitation scripts or rule-based logic and offering limited responsiveness to authentic classroom dynamics [11–13]. As a result, such systems may fail to capture learners' real needs, leading to shallow engagement or low adoption.

Educational research emphasises that the *quality* of participation, rather than the quantity of contributions, determines the educational value of collaborative discussion [14]. Productive discussion therefore requires explicit scaffolds that help learners share ideas, listen to peers, and deepen reasoning [12, 13]. Within the Academically Productive Talk (APT) framework [15], interactional talk moves—that is, brief facilitative prompts used to guide discussion—together with the Four Goals of Productive Discussions (FGPDs) [16] articulate key dimensions of high-quality dialogue. However, translating these abstract pedagogical principles into GenAI facilitation presents a significant design challenge. Unlike human facilitators, who can intuitively sense when to deploy specific talk moves (e.g., brief prompts such as "Can someone build on this idea?") GenAI systems require explicit computational logic to determine *when* and *why* to intervene during ongoing interaction.

Crucially, many breakdowns in authentic group discussion arise not from deficiencies in discourse moves alone, but from failures in shared regulation, such as coordinating participation, sustaining focus, and maintaining psychological safety. The framework of Socially Shared Regulation of Learning (SSRL) provides a lens for understanding how these cognitive, social, and emotional processes are jointly monitored and regulated over time [17, 18]. Yet, few existing AI-facilitated discussion systems explicitly model or support this regulatory dimension. Most current tools prioritise analysing epistemic content (what is said), rather than the regulatory conditions that sustain collaboration (how interaction unfolds) [19]. However, research on SSRL suggests that without regulating participation and focus, content-rich discussion alone is insufficient to sustain productive collaboration [18]. This leaves a critical gap in the design of GenAI facilitators for productive long-term group discussion.

To address these limitations, this study aims to design and validate an AI-facilitated approach for supporting productive collaborative discussions by operationalising theories of productive talk and socially shared regulation into explicit, theory-informed facilitation mechanisms. Grounded in the APT framework and its FGPDs [15, 16], this approach translates principles of productive talk into computationally actionable facilitation strategies. Insights from the SSRL framework further inform the design by shaping how AI facilitation can scaffold shared regulation processes—such as participation balance, pacing, and socio-emotional conditions—without undermining learners' epistemic agency. Adopting a design-based research (DBR) approach [20], the study conducts a formative analysis of learners' regulatory challenges, develops an AI-facilitated tool as a research instantiation, and evaluates its feasibility and perceived effectiveness in authentic collaborative settings.

This study makes three primary contributions. First, drawing on student interviews, we identify regulatory breakdowns in group discussion and derive learner-informed design requirements for GenAI facilitation. This addresses limited learner involvement in the early design of prior discussion tools [11, 12]. Second, we demonstrate how a low-interference GenAI facilitator can be designed and instantiated to support collaborative regulation, by grounding its modular supports in APT, FGPDs, and SSRL and showing how these abstractions can be operationalised in practice. Third, we report formative evidence of feasibility and perceived support and derive design guidelines for future GenAI-facilitated discussion tools.

Guided by these research objectives, the study addresses the following research questions:

- RQ1: What regulatory challenges do students experience in authentic group discussions, and what support do they expect from a GenAI facilitator?
- RQ2: How can a GenAI facilitator be designed to operationalise shared regulation principles to support productive discussion?
- RQ3: How does the prototype facilitate productive discussion across the FGPDs dimensions?

2 Related Work

2.1 Theoretical Foundations of Productive Discussion

Collaborative learning emphasises the quality of interaction—such as reasoning and elaboration—rather than simply the amount of participation [21]. The

APT framework addresses this by structuring dialogue through interactional talk moves(e.g., “Can someone build on this idea?”) [15]. The FGPDs characterise dialogue productivity by defining key behavioural and cognitive orientations in interaction. Specifically, they comprise four goal-level dimensions: (G1) sharing and clarifying ideas, (G2) orienting to and listening to others, (G3) deepening reasoning, and (G4) engaging with others’ reasoning [16]. Together, they specify *how* dialogue unfolds and *what* it aims to achieve. However, productive discourse often fails in authentic learning contexts due to “invisible” breakdowns like uneven participation or pacing issues [4, 5]. Discourse scaffolds alone are insufficient without regulating these underlying conditions. The SSRL framework conceptualises collaboration as a process of collective monitoring and regulation of cognitive, motivational, and socio-emotional states, helping identify *when* and *why* support is needed to sustain productive interaction [18, 17].

Effective facilitation thus requires a dual approach: SSRL informs when and why intervention is needed (regulation), while APT/FGPDs guide how to enact it (discourse). Accordingly, this study addresses the identified gap by designing and evaluating a GenAI-facilitated approach that integrates regulatory logic with discourse-level scaffolding to support productive group discussion.

2.2 AI Facilitation in CSCL

Early CSCL tools focused on coordination but lacked contingent dialogic support—support that adapts to changes in group interaction (e.g., silence or uneven participation) [22]. To address this, conversational agents were introduced. For instance, *MentorChat* [8] embedded scripted APT-based prompts to encourage elaboration, while *MOCA* [7] employed motivational language to enhance social presence. Although promising, these rule-based interventions often lacked responsiveness to emergent breakdowns like sustained silence. More recent systems, such as *Clair* [12], combine learning analytics with real-time dialogue analysis. While improving behavioural responsiveness, they primarily focus on epistemic content rather than regulating socio-emotional conditions. From an SSRL perspective, this limits their ability to sustain productive collaboration over time, as socio-emotional breakdowns often precede epistemic difficulties [18].

GenAI expands design possibilities through context-aware generation. However, learners often perceive GenAI tools as *authoritative information sources* or passive answer providers, leading to reduced agency and over-assistance [5, 23]. While group awareness tools(GATs), such as the PAT [24], increase transparency by visualising participation [25], they rarely enact active interventions.

Taken together, prior work has largely treated discourse support and regulatory support in isolation, leaving a gap in integrating discourse frameworks (APT/FGPDs) with regulatory logic (SSRL) for contingent facilitation. Recent reviews further highlight the need for systems that combine behavioural responsiveness with theory-grounded scaffolds [26, 27]. Addressing this gap calls for learner-informed, theory-grounded designs that derive design requirements from students and involve them in the design and validation phases.

3 Phase 1: Requirements Elicitation

3.1 Method

A semi-structured interview approach was adopted [28, 29]. This method enabled participants to articulate individual and contextualised experiences of group discussion, while maintaining coherence across sessions and flexibility to probe emergent issues [30, 31]. Fifteen IT students (n=15) voluntarily participated in this phase. All had prior experience with collaborative coursework or group projects, ensuring relevance to the study context. Before the interviews, participants completed a short questionnaire capturing background information (e.g., discipline and study stage), prior experience with GenAI tools, and details of their most recent group discussion or project.

Each interview lasted approximately 30–45 minutes and was conducted either online (i.e., Zoom) or in person on campus. The interview guide was organised around five themes (the full protocol is provided in the online supplementary materials[1]): (1) students' recent experiences of group collaboration; (2) their conceptions of productive discussion and its indicators (e.g., reasoning with evidence, attentive listening, consensus building); (3) barriers and breakdown patterns in discussion (e.g., silence, dominance, topic drift, time pressure); (4) existing support mechanisms and their perceived effectiveness; and (5) expectations regarding the roles of GenAI facilitation. All interviews were audio-recorded with informed consent, transcribed verbatim, and anonymised (P1–P15). Quotes were lightly edited for readability without altering meaning.

3.2 Analysis

A thematic analysis [32] was conducted in three stages. **Stage 1: Inductive open coding** involved line-by-line reading to summarise individual discussion experiences and challenges, specifically focusing on interaction breakdowns (e.g., silence, dominance). The coding scheme is provided in the online supplementary materials[1]. **Stage 2: Data-to-theory alignment** interpreted these recurring patterns through the FGPDs and SSRL frameworks, mapping empirical challenges to breakdowns in the four productive discussion goals (sharing, listening, reasoning, consensus). **Stage 3: Expert validation** strengthened the credibility of the analysis through independent review by two experts in educational technology and learning analytics, who examined the identified themes, their theoretical interpretations, and representative excerpts. Any discrepancies regarding theme definitions were resolved through discussion, resulting in full consensus on the final set of five functional themes reported in the findings.

3.3 Findings

Across interviews, students consistently described recurring *regulatory breakdowns* in participation, pacing, and collective monitoring that constrained pro-

[1] https://github.com/plsybla52115/ProDAIS-AIED2026

ductive discussion. These patterns reflected not isolated communication problems but broader difficulties in sustaining socially shared regulation during group works. As summarised in Table 1, each observed breakdown was diagnosed through SSRL (why it occurred) and linked to the FGPDs it undermined (what to support), which together guided the system's design focus. Students also emphasised that AI support should remain lightweight and non-intrusive.

Structured Agenda and Pacing Support. Students consistently emphasised that well-defined agendas were essential to maintaining efficiency and direction in group discussions. However, 60% of participants (n=9) reported that, without explicit time cues, their discussions often drifted away from the main goal: *"Sometimes we just talked about random things and forgot the main purpose."* (P3). These accounts indicated breakdowns in shared temporal regulation, where groups struggled to collectively monitor progress and manage transitions between discussion stages. In response, students expressed a preference for low-interference AI facilitation that could gently support shared temporal regulation by providing unobtrusive reminders to sustain pacing and stage progression.

Equitable Participation and Quiet-Member Support. Uneven participation was the most frequently reported challenge (n=12, 80%). Students described scenarios in which dominant speakers shaped the flow of discussion while quieter members gradually withdrew: *"When one person dominates, others tend to disengage."* (P7). Such patterns indicated difficulties in regulating participation equity and sustaining the socio-emotional conditions needed for inclusive discussion. In response, students expressed a preference for gentle and non-evaluative AI facilitation that encourages quieter members to contribute without creating pressure or anxiety, thereby supporting balanced participation while preserving learner agency.

Shared Memory and Summarisation. Participants identified the absence of a shared discussion record as a major barrier to continuity and follow-up (n=10, 67%). As P9 explained, *"If no one took notes, we often forgot decisions by the following week."* This challenge indicated a breakdown in collective monitoring and shared understanding across discussion episodes. Consequently, students expected AI facilitation to provide periodic and post-meeting summaries that externalise key ideas and decisions, supporting collective memory and helping groups maintain a shared sense of progress over time.

Action Orientation and Accountability. Discussions often failed to translate into concrete outcomes due to unclear responsibility allocation (n=8, 53%). Students recalled instances where work distribution became uneven, with a small number of members taking on most tasks: *"Two or three people ended up doing everything."* (P8). These experiences highlighted weaknesses in regulating accountability and follow-up actions after discussion. As a result, students hoped that AI facilitation would help consolidate decisions into actionable to-do items, supporting clearer responsibility sharing without taking over decision-making.

Re-activation and Brainstorming Support. When discussions stalled or fell into silence (n=7, 47%), students expressed a desire for lightweight prompt to help re-engage collective thinking. As one participant suggested, *"AI could give*

Table 1. Mapping Observed Breakdowns to Theory-Informed Design Focus

Observed Breakdown	**FGPDs Goal**	**SSRL Interpretation**	**Design Focus**
Uneven participation & prolonged silence	Listening; Sharing	Failure to sustain shared regulation of participation and socio-emotional safety.	Support Quiet
Topic drift & weak pacing	Engaging; Reasoning	Insufficient regulation of task focus and temporal progress.	Agenda (Action items)
Loss of ideas between meetings	Sharing; Engaging	Breakdown in collective monitoring and shared memory.	STT & Summarisation
Unclear responsibility	Sharing	Weak accountability regulation and follow-up coordination.	Agenda & Summarisation
Stagnation / Lack of ideas	Reasoning	Difficulty re-activating shared cognitive engagement.	Brainstorming

us 3–4 new ideas when the room goes quiet." (P7). Importantly, students did not expect the AI to control the discussion. Instead, they envisioned a facilitative role in which GenAI provides timely prompts that help reactivate shared cognitive engagement while preserving learners' ownership of ideas and decisions.

Together, these findings outlined empirically grounded requirements for GenAI facilitation that supports productive discussion by strengthening shared regulation of participation, pacing, memory, and engagement. These requirements directly informed the design of the five core modules of the Productive Discussion AI System (ProDAIS) described in Phase 2.

4 Phase 2: System Prototype Design and Evaluation

4.1 Design Rationale and Architecture

Informed by the regulatory breakdowns identified in Phase 1, ProDAIS was designed as a low-interference *GenAI facilitator*. Rather than directing content, the system supported the socio-cognitive *conditions* for productive interaction—such as equity, focus, and shared memory—while preserving learner agency.

The design operationalised a tripartite theoretical grounding: FGPDs specified *what* dimensions to support (sharing, listening, reasoning, and engaging with others' reasoning); APT informed *how* interventions were phrased (e.g., non-evaluative inviting); and SSRL guided *when* and *why* to intervene based on detected breakdowns. From a system perspective, ProDAIS followed a cyclical workflow: *Sensing* (monitoring speech/time) → *Interpretation* (detecting regulatory signals) → *Intervention*, with timing governed by a minimal interruption logic [33] to ensure that support enhances rather than disrupts the flow of discussion. Figure 1 illustrates this architecture.

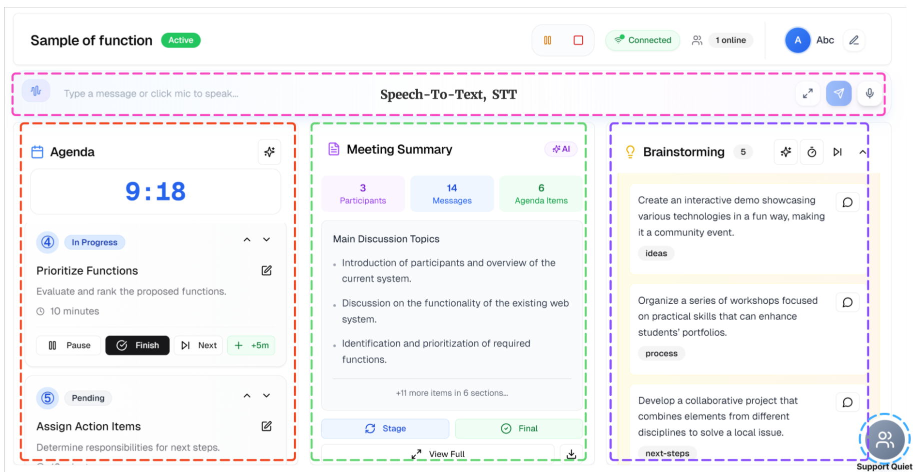

Fig. 1. ProDAIS system architecture and core module overview

4.2 Core Functional Modules of ProDAIS

Each module was designed to address specific needs identified in Phase 1 by translating theoretical principles into concrete computational functions:

Speech-to-Text (STT). The STT module transformed spoken contributions into a shared textual record that served as a collective memory. By externalising talk into visible text, the transcript functioned as a regulatory artefact supporting collective monitoring and grounding, enabling groups and other AI modules to revisit prior ideas when coordinating the discussion.

Summarisation. To address "topic drift" and fragmented understanding, the Summarisation module implemented progressive summarisation. At regular intervals or on demand, the system generated concise summaries highlighting current goals, key ideas, and action points, helping groups review progress and maintain shared understanding during discussion.

Support Quiet. This module targeted participation imbalance by detecting prolonged silence and interaction patterns. When sustained imbalance was identified, the system generated brief, non-intrusive "talk move" prompts (e.g., gentle invitations), framed at the group level. This design regulated participation socially, reducing interpersonal risk and fostering more equitable contribution.

Agenda. The Agenda module externalised temporal regulation through visual timeboxing and stage-based reminders. By signalling approaching transitions and allowing flexible adjustments, the module made pacing visible to the group, supporting shared time awareness and coordination during discussion.

Table 2. Perceived support for FGPDs dimensions ($N = 15$, 1–5 Scale)

FGPDs	Mean (SD)	Representative User Feedback
Sharing	4.24 (0.37)	"The summary helped us make vague ideas clearer." (P6)
Listening	4.49 (0.35)	"Inviting quieter members helped the group pay attention to different voices." (P8)
Reasoning	3.71 (0.72)	"Brainstorming gave us new directions to think about." (P3)
Engaging	4.47 (0.52)	"The agenda timer helped us stay on task." (P11)

Brainstorming. When interaction stalled or idea generation declined, this module offered lightweight, categorised prompts to stimulate divergent thinking. They provided directional guidance without supplying solutions, helping groups explore alternative ideas and maintain momentum during discussion.

Together, these modules enabled ProDAIS to support group discussion through minimal, theory-informed interventions. Detailed interaction logic and implementation specifics are provided in the online supplementary materials[2].

4.3 Evaluation

To assess the efficacy of ProDAIS as a regulatory facilitator, a user study was conducted focusing on facilitation effectiveness grounded in the FGPDs [16], usability measured using the SUS [34], and perceived cognitive load assessed via selected dimensions of the NASA Task Load Index (NASA-TLX) [35].

Participants and Task. Fifteen IT students (N=15, 5 groups) took part in the study, all with prior collaborative experience. All groups completed the same 45-minute task of co-designing a paper prototype for a "Healthy Lifestyle App", which required negotiation and consensus building.

Procedure and Data Collection. Following a 10-minute tutorial, groups completed the task using ProDAIS. Data were collected via: (1) **system logs** (intervention timestamps); (2) **surveys** including SUS, NASA-TLX (1–7 scale), and a custom FGPDs-based instrument (1–5 Likert scale); and (3) **open-ended responses** for qualitative feedback.

4.4 Findings

Facilitation Effectiveness Quantitative results (Table 2) showed positive ratings across all four FGPDs dimensions, with relatively small standard deviations.

The highest ratings were reported for **Listening** (M=4.49, SD=0.35) and **Engaging** (M=4.47, SD=0.52). Module-level items showed a similar pattern: the *Agenda* module received the highest specific praise (M=4.73, SD=0.59 for "maintaining progress"), and was frequently described as helpful for keeping the group in synchronisation. The *Summarisation* module (M=4.67, SD=0.62 for

[2] https://github.com/plsybla52115/ProDAIS-AIED2026

"helping understanding") was perceived as useful for reducing the "topic drift" identified in Phase 1. Although *Reasoning* received slightly lower ratings overall (M=3.71, SD=0.72), students still viewed the *Brainstorming* module as helpful for generating new ideas (M=3.86, SD=0.86). Several participants commented that, while the prompts were effective in breaking silence, their perceived usefulness varied across discussion contexts.

Usability and Cognitive Load. The system achieved a mean SUS score of 78.50 (SD=14.90, Median=80.00), which falls within the commonly reported "Good-to-Excellent" usability range. Regarding cognitive load, participants reported high perceived *Performance* (M=5.67, SD=1.18) and notably low *Frustration* (M=1.67, SD=0.82). *Mental Demand* was moderate (M=3.60, SD=1.84) and *Effort* was reasonable (M=4.60, SD=1.35), suggesting that, in students' view, the system helped with regulatory tasks (e.g., timekeeping, note-taking) without introducing substantial additional cognitive burden.

Perceived Usefulness. Participants ranked the modules (1=Best, 5=Worst) in the following order of preference: **Summarisation** ($M = 1.87, SD = 1.13$) was considered the most useful feature, followed by **STT** (M=2.07, SD=1.03) and **Agenda** (M=2.07, SD=1.03). Proactive interventions such as **Support Quiet** (M=2.80, SD=1.66) and **Brainstorming** ($M = 3.07, SD = 1.53$) were ranked slightly lower. This preference pattern suggests that students tended to value structural scaffolds for memory and time management over more direct social or cognitive interventions, aligning with a low-interference facilitation approach.

5 Discussion

This section synthesises findings across two phases to examine how GenAI can function as a low-interference facilitator for productive group discussion, and to derive design insights for future GenAI-facilitated discussion tools.

5.1 Regulatory breakdowns and value of externalising coordinations

Our findings indicated that the most persistent challenges to productive group discussion were primarily regulatory rather than epistemic. Across both phases, participants rarely reported a lack of ideas; instead, they struggled to sustain balanced participation, pacing, and shared focus over time. This pattern is reflected in Phase 2, where structural modules (Agenda, Summarisation) were ranked as more useful than generative modules (e.g., Brainstorming). The lower rating for reasoning suggests that while lightweight GenAI prompts successfully stimulated divergent thinking (generating ideas), they were less effective in supporting convergent thinking (integrating and deepening ideas), which likely requires more context-aware scaffolding. This aligns with work on SSRL, which argues that productive collaboration is not only a matter of generating contributions but of jointly monitoring, coordinating, and refining them over time [17, 18].

In this context, learners valued support that externalised regulation through lightweight shared artefacts. By converting interactional elements (e.g., contributions, timing, decisions) into shared and persistent resources, ProDAIS reduced the coordination overhead of tracking participation and progress. The high perceived utility of the Summarisation module, for instance, suggests that students rely on AI not to do the thinking for them, but to maintain the "shared memory" required for collective reasoning. This aligns with SSRL perspectives that characterise effective collaboration as a process of distributed monitoring and control that is made visible and jointly accessible at the group level [36].

5.2 Low-interference facilitation and the tension between regulation and intrusion

Our findings delineate important boundaries in how AI-driven regulatory support is experienced. While externalising regulatory states (e.g., via the Agenda) can support coordination, proactive prompts that explicitly surfaced silence or participation imbalance (e.g., the Support Quiet module) were sometimes perceived as intrusive or less useful. This is consistent with findings in human–AI collaboration research showing that over-explicit or poorly timed interventions can disrupt social dynamics and undermine psychological safety [37].

These reactions indicate that the perceived helpfulness of GenAI facilitation depends not only on what is supported, but critically on how and when support is introduced [38, 39]. Interventions that regulate the conditions for talk—such as participation equity, pacing cues, and shared grounding—were more readily accepted when framed as brief, invitation-based prompts or shared artefacts, rather than as direct signals of imbalance.

This finding highlights a key design tension between regulation and intrusion. Structural scaffolds that sustain the prerequisites for productive talk can enable sharing, listening, and consensus-building without appropriating learners' epistemic work. However, when regulatory support becomes overly explicit or socially exposing, it risks undermining the very conditions it seeks to support. This underscores the importance of low-interference, context-sensitive facilitation strategies that preserve learner agency and psychological safety while still making coordination visible.

5.3 Implications for Research and Practice

For research, these findings suggest that GenAI facilitation in collaborative discussion should be studied as regulatory support embedded in collaboration over time, rather than as isolated prompt or content design. This aligns with recent work that conceptualises collaboration as a dynamic coordination process and emphasises the role of timing, visibility, and perceived intrusiveness in human–AI collaboration [33, 38]. From this viewpoint, modelling regulatory breakdowns and examining learners' responses to different forms of AI-supported regulation constitute a promising analytical lens for understanding human–AI collaborative facilitation.

For practice, three design guidelines for AI-facilitated collaboration tools are suggested: (i) prioritise support for regulatory conditions rather than content generation; (ii) externalise regulation through shared and visible artefacts; and (iii) design interventions to be infrequent, lightweight, and invitation-based, minimising disruption to learner agency and ongoing group dynamics.

5.4 Limitations and Future Work

This study has several limitations. The evaluation was conducted with a small number of short-term collaborative sessions, and participants may have been more open to experimenting with AI than typical classroom learners. In addition, the current prototype has limited capacity to interpret nuanced social cues, affective states, and contextual subtleties that shape real-world discussion dynamics. Accordingly, the findings should be interpreted as formative and design-oriented rather than as evidence of learning effectiveness.

Future work will extend this research in three directions. First, ProDAIS will be deployed in authentic classroom settings to examine its role in longer-term collaboration and learning processes. Second, future iterations will explore adaptive and affect-aware facilitation to better align AI interventions with group dynamics and socio-emotional conditions. Third, the design will be expanded toward educator-facing analytics that support teachers in observing and coordinating group processes alongside AI facilitation.

6 Conclusion

This study examined how Generative AI can facilitate productive group discussion by addressing regulatory breakdowns that often hinder collaboration. Grounded in the SSRL framework and informed by APT and FGPDs, we designed ProDAIS as a low-interference GenAI facilitator that supports time management, balanced participation, and shared focus during group discussion.

The findings highlight a shift in how AI support for CSCL can be conceptualised. Rather than acting as a content provider or "knowledge oracle", GenAI can function as a **regulatory facilitator** that supports the conditions for productive dialogue while leaving meaning-making to learners. By externalising coordination processes and intervening only when breakdowns occur, AI support can augment learner agency without disrupting ongoing interaction.

References

1. Gerry Stahl. *Group cognition: Computer support for building collaborative knowledge*. The MIT Press, 2006.
2. Pierre Dillenbourg. What do you mean by collaborative learning? *Collaborative-learning: Cognitive and computational approaches.*, pages 1–19, 1999.
3. Elisabeth G Cohen and Rachel A Lotan. *Designing groupwork: strategies for the heterogeneous classroom third edition*. Teachers College Press, 2014.

4. Genimon Vadakkemulanjanal Joseph, P Athira, M Anit Thomas, Dawn Jose, Therese V Roy, and Malavika Prasad. Impact of digital literacy, use of ai tools and peer collaboration on ai assisted learning: Perceptions of the university students. *Digital Education Review*, 45:43–49, 2024.
5. Guangrui Fan, Dandan Liu, Rui Zhang, and Lihu Pan. The impact of ai-assisted pair programming on student motivation, programming anxiety, collaborative learning, and programming performance: a comparative study with traditional pair programming and individual approaches. *IJSTEM Education*, 12(1):16, 2025.
6. Tong Wu, Xiaohang Tang, Sam Wong, Xi Chen, Clifford A Shaffer, and Yan Chen. The impact of group discussion and formation on student performance: An experience report in a large cs1 course. In *ACM SIGCSE*, pages 1260–1266, 2025.
7. Tao Xie, Ruobin Liu, Yijin Chen, and Geping Liu. Moca: A motivational online conversational agent for improving student engagement in collaborative learning. *IEEE Transactions on Learning Technologies*, 14(5):653–664, 2021.
8. Stergios Tegos, Stavros Demetriadis, and Anastasios Karakostas. Promoting academically productive talk with conversational agent interventions in collaborative learning settings. *Computers & Education*, 87:309–325, 2015.
9. Kok-Sing Tang and Gde Buana Sandila Putra. Generative ai as a dialogic partner: Enhancing multiple perspectives, reasoning, and argumentation in science education with customized chatbots. *JSET*, pages 1–13, 2025.
10. Pegah Karimi, Jeba Rezwana, Safat Siddiqui, Mary Lou Maher, and Nasrin Dehbozorgi. Creative sketching partner: an analysis of human-ai co-creativity. In *25th IUI*, pages 221–230, 2020.
11. Riordan Alfredo, Vanessa Echeverria, Yueqiao Jin, Lixiang Yan, Zachari Swiecki, Dragan Gašević, and Roberto Martinez-Maldonado. Human-centred learning analytics and ai in education: A systematic literature review. *Computers and Education: Artificial Intelligence*, 6:100215, 2024.
12. Adelson de Araujo, Pantelis M Papadopoulos, Susan McKenney, and Ton de Jong. A learning analytics-based collaborative conversational agent to foster productive dialogue in inquiry learning. *JCAL*, 40(6):2700–2714, 2024.
13. Stergios Tegos, Stavros Demetriadis, Pantelis M Papadopoulos, and Armin Weinberger. Conversational agents for academically productive talk: A comparison of directed and undirected agent interventions. *ijCSCL*, 11:417–440, 2016.
14. Neil Mercer. *Words and minds: How we use language to think together*. Routledge, 2002.
15. Sarah Michaels, MC O'Connor, MW Hall, and L Resnick. Accountable talk: Classroom conversation that works. *Pittsburg: University of Pittsburgh*, 2002.
16. Sarah Michaels and Catherine O'Connor. Conceptualizing talk moves as tools: Professional development approaches for academically productive discussion. *Socializing intelligence through talk and dialogue*, pages 347–362, 2015.
17. Allyson Hadwin and Mika Oshige. Self-regulation, coregulation, and socially shared regulation: Exploring perspectives of social in self-regulated learning theory. *Teachers College Record*, 113(2):240–264, 2011.
18. Sanna Järvelä, Paul A Kirschner, Ernesto Panadero, Jonna Malmberg, Chris Phielix, Jos Jaspers, Marika Koivuniemi, and Hanna Järvenoja. Enhancing socially shared regulation in collaborative learning groups: Designing for cscl regulation tools. *Educational Technology R&D*, 63(1):125–142, 2015.
19. Sanna Järvelä, Andy Nguyen, and Allyson Hadwin. Human and artificial intelligence collaboration for socially shared regulation in learning. *British Journal of Educational Technology*, 54(5):1057–1076, 2023.

20. Feng Wang and Michael J Hannafin. Design-based research and technology-enhanced learning environments. *ETRD*, 53(4):5–23, 2005.
21. Michelene Chi and Ruth Wylie. The icap framework: Linking cognitive engagement to active learning outcomes. *Educational psychologist*, 49(4):219–243, 2014.
22. Qiaosi Wang, Ida Camacho, and Ashok K Goel. Investigating the potential of ai-based social matching systems to facilitate social interaction among online learners. pages 279–298. Springer, 2022.
23. Lasha Labadze, Maya Grigolia, and Lela Machaidze. Role of ai chatbots in education: systematic literature review. *ijETHE*, 20(1):56, 2023.
24. Marie Guillon and Gaëlle Molinari. A group awareness tool for self-assessment and visualization of participation: Its effect on the regulation of unequal participation. In *ijCSCL 2023, pp. 35-42*. International Society of the Learning Sciences, 2023.
25. Jeroen Janssen, Gijsbert Erkens, and Paul A Kirschner. Group awareness tools: It's what you do with it that matters. *Computers in Human Behavior*, 27(3):1046–1058, 2011.
26. Liru Hu and Gaowei Chen. A systematic review and meta-analysis of productive peer talk moves. *Journal of Behavioral Education*, 33(4):798–830, 2024.
27. Zhanxin Hao, Jianxiao Jiang, Jifan Yu, Zhiyuan Liu, and Yu Zhang. Student engagement in collaborative learning with ai agents in an llm-empowered learning environment: A cluster analysis. *arXiv preprint arXiv:2503.01694*, 2025.
28. Rosalind Edwards and Janet Holland. *What is qualitative interviewing?* Bloomsbury Academic, 2013.
29. Nigel King. *The qualitative research interview.* Sage Publications, Inc, 1994.
30. Nigel King, Joanna Brooks, and Christine Horrocks. *Interviews in qualitative research.* SAGE Publications ltd, 2018.
31. Steinar Kvale and Svend Brinkmann. *Interviews: Learning the craft of qualitative research interviewing.* sage, 2009.
32. Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.
33. Yang Deng, Lizi Liao, Zhonghua Zheng, Grace Hui Yang, and Tat-Seng Chua. Towards human-centered proactive conversational agents. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 807–818, 2024.
34. John Brooke et al. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.
35. Sandra G Hart. Nasa-task load index (nasa-tlx); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 50, pages 904–908. Sage publications Sage CA: Los Angeles, CA, 2006.
36. Y. Yang et al. Development and validation of the socially shared regulation of learning (ssrl) questionnaire. *Frontiers in Psychology*, 2025.
37. Amy C Edmondson and Zhike Lei. Psychological safety: The history, renaissance, and future of an interpersonal construct. *Annu. Rev. Organ. Psychol. Organ. Behav.*, 1(1):23–43, 2014.
38. Thomas Schmutz, Lawrence J. Hettinger, et al. Collaboration between humans and artificial intelligence. *Current Opinion in Psychology*, 59:101781, 2024.
39. Sander de Jong, Joel Wester, Tim Schrills, Kristina S. Secher, Carla F. Griggio, and Niels van Berkel. Assessing cognitive and social awareness among group members in ai-assisted collaboration. In *MUM '24*, pages 338–350, 2024.