DiscourseMap: Enhancing Epistemic Visibility in Asynchronous Discussions

Jiin Hur

Teachers College, Columbia University

DiscourseMap: Enhancing Epistemic Visibility in Asynchronous Discussion

In graduate education, asynchronous discussion forums are designed to support reflective and student-led dialogue. However, in reality, students tend to post minimally, choose reply targets based on convenience, and stop participating once they fulfill the participation requirements. This pattern is not primarily a problem of lack of motivation. Typical learning management system (LMS) forums present posts as flat threads in chronological order, making it difficult to grasp how ideas are connected, where perspectives diverge, and where meaningful contributions can add value. When this visibility is lacking, students only participate superficially, and instructors find it hard to guide deep discussions.

DiscourseMap makes the epistemic landscape of a discussion visible in real time in order to address this structural limitation. Implemented as a Canvas plugin, this system transforms linear discussion threads into a dynamic spatial map that represents thematic clusters and their relationships, so that learners can identify how ideas are connected, where discussions are concentrated, and which areas have not yet been sufficiently explored. In addition, this system highlights gaps, weak connections, and dominant perspectives, and provides context-sensitive prompts to extend, challenge, and connect ideas. Grounded in four theoretical frameworks including knowledge building, group awareness tools, multimodal visualization, and AI-based scaffolding, this project explores whether making the structure of discourse visible can support more purpose-oriented participation and promote knowledge-building-centered interaction.

Design Rationale

Theoretical Framework

Structural Limitations of Asynchronous Discussion

Higher education treats discussion as a core method of discourse-based learning in which learners adjust their own reasoning, examine diverse perspectives, and deepen their understanding of complex issues (O'Brien, 1939; Parker & Hess, 2001; Abdulbaki et al., 2018; Ying, 2020). As online learning expands, LMS forums have become the primary discussion space (Yoshida et al., 2026). However, a large body of prior research points out a persistent gap between these educational ideals and actual participation patterns. This gap appears as fragmented interaction with parallel monologues and agreement-driven responses, which limits the development of more substantive and exploratory dialogue (Cal Varela & Fernández Polo, 2022). When responses cluster around initial posts, dominant interpretations form and discourage challenge (Wise et al., 2014). In addition, when participation rules and evaluation criteria are rigid, learners prioritize fulfilling requirements rather than expanding inquiry (Kipling et al., 2023). As a result, commonly used indicators such as the number of posts or the frequency of comments do not sufficiently reflect the quality of participation or the depth of learning in such environments (Gasell et al., 2022).

Knowledge Building and LMS Discussion

Discussion requires more than mere participation or information exchange; it must advance the group's shared understanding. The Knowledge Building (KB) theory of Scardamalia and Bereiter (2006) explains such discourse as a process oriented toward the advancement of collective understanding and the expansion of the knowledge base. However, typical LMS discussion environments, (e.g., Canvas Discussion Forum), operate centered on formal participation, such as the quantity of posts or responses, and therefore tend to remain at the level of knowledge sharing that expresses existing understanding rather than supporting KB (Cal-Varela & Fernández-Polo, 2022; Gasell et al., 2022; Kipling et al., 2023).

This gap between KB theory and LMS practice is closely related to the structural characteristics of discussion boards. Threaded structures obscure relationships among ideas and across contributions, and as a result, discourse becomes fragmented, and cumulative advancement becomes difficult to achieve. For example, Van Aalst (2009) empirically confirms this problem and explains that actual online discourse tends to remain at the level of knowledge sharing. He points out that the quality of discourse is determined not by the quantity of participation but by the nature of contributions, and considers activities such as explanation-seeking questions, the interpretation of information, and the connection among concepts to be essential. However, because these differences are not made visible in typical discussion environments, learners may find it difficult to understand what meaning their contributions have within the discourse. In addition, knowledge-building discourse is closely connected to social interaction, and socio-cognitive conflict that arises from the confrontation between different perspectives drives deeper understanding (Long et al., 2020). However, sequential thread structures weaken these interactions and, in practice, may reinforce participation centered on surface-level responses (Asterhan, 2015).

Taken together, this suggests that knowledge-building discourse cannot be achieved through a mere increase in participation, but requires structural conditions that make visible the relationships among ideas, the meaning of contributions, and the dynamics of interaction. However, such conditions are largely absent in typical LMS discussion environments, where learners have limited access to information about others' understanding, perspectives, and contributions. As a result, learners tend to participate based on an incomplete picture of the group's epistemic state, which can reinforce surface-level agreement and constrain the kind of productive conflict and idea-building that deeper discourse requires.

Group Awareness Tool (GAT) as Structural Support

Researchers propose Group Awareness Tools (GAT) as a structural response to this lack of mutual epistemic awareness. GAT is the technology that supports collaborative interaction by making learners' cognitive states and participation patterns visible to other participants (Bodemer et al., 2018). In existing research, this is categorized into social GAT and cognitive GAT, where social GAT provides information such as participation frequency or activity level, whereas cognitive GAT reveals learners' knowledge states, perspectives, and uncertainties. Cognitive GAT enables more meaningful interaction beyond simple agreement by making visible the differences in understanding and perspectives among learners. When different interpretations or knowledge states are revealed, learners can recognize the points where discussion is needed and participate by requesting explanations or reconstructing existing ideas (Bodemer et al., 2018). This process enables socio-cognitive conflict, which deepens understanding (Long et al., 2020). To support this kind of process, researchers have also explored how making the discourse itself more transparent through visualization can encourage deeper engagement. Kauer et al. (2024) show that visualization that spatially aggregates contributions can reveal implicit knowledge and participation structures, and Asterhan (2015) emphasizes that deeper participation becomes possible when learners are able to recognize the epistemic work taking place in the discourse. The design of collaborative discussion environments must go beyond simply showing the level of participation and should be able to reveal epistemic structures such as differences in perspectives, uncertainties, and the distribution of knowledge.

Multimodal Visualization and AI-mediated Scaffolding

Another important problem that appears in online discussion environments is that, as discourse information is presented in a text-centered manner, interpreting discourse imposes high cognitive load. The linear and accumulative structure of posts makes it difficult to grasp the flow of discussion and the key issues, and this can function as a factor that keeps learners' participation at a surface level. To complement these limitations, multimodal representations that visually reconstruct discourse information are proposed. Rusen (2025) points out that multimodal representations support discourse comprehension and engagement, and particularly emphasizes that such representational forms are key elements that support discourse understanding in online environments. Jang and Lee (2024) present the main topics of discussion in a structured manner by utilizing topic modeling-based network visualization and show that such visualization is effective in promoting learners' reflective thinking. In particular, higher usability of the visualization tools leads to more meaningful interpretation.

Along with this, recent research shows that AI can play an important mediating role in such multimodal environments. Dovhaniuk et al. (2025) explain that AI acts as a semantic mediator. It restructures discourse and supports interpretation. In this context, AI restructures discourse into interpretable forms. In addition, AI feedback and scaffolding improve participation and learning outcomes. Ba et al. (2025) show that immediate and personalized feedback provided by AI can enhance learners' level of inquiry and induce more balanced participation. These effects appear particularly when alleviating the cognitive burden that learners experience in the process of generating their contributions. Meanwhile, Shea et al. (2006) emphasize that the instructor's role in discourse design and facilitation is central to the formation of a learning community in online learning, which suggests that AI can function in a way that complements such teaching presence. Multimodal visualization and AI-mediated scaffolding perform key roles in making the structure of discourse understandable, reducing learners' cognitive burden, and supporting deeper participation. This suggests that in order to improve the quality of discourse in online discussion environments, an integrated approach is necessary.

Population and Context

DiscourseMap is designed for graduate-level classes or seminars, specifically those centered around the critical analysis of multiple professor-selected materials. Given that the humanities and social sciences rely heavily on discursive participation (Reed & Alexander, 2009), these fields represent the primary application for the platform.

DiscourseMap acts as an AI assistant for synthesizing discussion. Students interact with a structured knowledge map instead of navigating cluttered, linear threads. As the system automatically identifies and tags contributions, which are distinguishing between "evidence," "limitations," and "claims," learners gain visibility into the logical flow of the discussion. This scaffolding shifts students from isolated reflections to relational thinking.

Learning Goals and Structure

DiscourseMap's primary goal is the enhancement of epistemic performance within graduate-level discussions. Based on the framework of Barzilai and Chinn (2018), who identified key aspects of epistemic education, the learning goals of DiscourseMap, aligned with this framework, are structured as follows:

First, students "engage in reliable cognitive processes that lead to the achievement of epistemic aims" (Barzilai & Chinn, 2018). The system functions as an epistemic GAT, specifically designed to surface the epistemic landscape of a conversation. By making diverging perspectives visible (Bodemer et al., 2018) and providing clear "epistemic guidance" (Asterhan, 2015), DiscourseMap renders the collective structure of engagement explicit through "discursive patinas"— the spatial aggregation of contributions anchored to a shared visualization (Kauer et al., 2024; Zhu et al., 2026). This visibility prompts learners to revise their understanding when encountering conflicting perspectives (Long et al., 2020). Furthermore, to reduce cognitive load, the tool integrates text with topic modeling-based network visualizations (Jang & Lee, 2024). In this multimodal environment, AI acts as a "semiotic actor," helping students navigate argumentation pathways and making discourse data more meaningful for reflection (Dovhaniuk et al., 2025).

Second, learners "participate in epistemic performance together with others" (Barzilai & Chinn, 2018). One core objective of DiscourseMap is to facilitate collaborative KB, moving students beyond mere knowledge-telling (Scardamalia & Bereiter, 2006). By enabling shared agency (Dovhaniuk et al., 2025), the system supports intersubjective meaning-making via shared artifacts (Stahl & Hakkarainen, 2021). This fosters knowledge synthesis, integrating fragmented ideas to generate conceptual innovation (Shui & Zhu, 2026).

Third, through interaction with DiscourseMap, learners "metacognitively regulate and understand epistemic performance" while adapting it to diverse situations (Barzilai & Chinn, 2018). AI-driven scaffolding supplements teaching presence by providing the facilitation traditionally handled by instructors (Shea et al., 2006). By providing personalized, context-sensitive prompts based on system-generated epistemic maps, the tool reduces the cognitive burden of generating high-quality responses and supports more balanced participation (Ba et al., 2025; Rusen, 2025). Ultimately, DiscourseMap supports metacognitive monitoring and strategic participation, fostering a high level of agency in diverse KB environments.

Assessment Strategy

To capture the growth of epistemic performance through DiscourseMap from multiple angles, the KB framework can be employed, as the system prioritizes the evolution of community knowledge. Accordingly, the assessment strategy can integrate instructor-led analytic rubrics with learner reflection journals to capture epistemic growth from multiple perspectives. The potential of rubrics to identify specific areas for improvement in academic performance has been well-proved (Reddy & Andrade, 2010). Simultaneously, reflective journals serve as a critical tool for developing metacognitive awareness and refining academic writing skills

(O'Connell & Dyment, 2006). This complementary approach captures both objective performance and learner reflection.

First, to evaluate the first goal (Reliable Cognitive Processes), the assessment strategy compares a student's initial post with their subsequent revisions following the 'thematic cluster' visualization. Within the KB framework, all ideas are treated as improvable objects that must be continuously refined for higher-level understanding (Scardamalia & Bereiter, 2006). Evaluation examines how students integrate cluster structures into their arguments. Supplementally, system logs assess whether students use the map as a cognitive resource.

Second, regarding the second goal (Participating Together), the assessment strategy analyzes the 'Connectivity' between posts. This involves measuring the frequency with which students use peers' ideas as materials for knowledge synthesis. In KB, knowledge is viewed as a 'public good' managed collectively (Scardamalia & Bereiter, 2006). Evaluation rubrics include the reduction of isolated nodes and the statistical increase in highly connected nodes, where conceptual innovations emerge from combining multiple perspectives, compared to simple responses.

Third, to assess the third goal (Metacognitive Regulation & Adaptation), the assessment strategy monitors whether learners modify their inquiry strategies following interaction with AI partners or visual maps subsequently. This aligns with 'Epistemic Agency', the capacity of learners to self-direct their knowledge-building journey by identifying what they do and do not know (Scardamalia & Bereiter, 2006). We measure the frequency of student interventions in 'logical collision points' or 'empty spaces' on the map. This evaluation determines how effectively AI 'semiotic actors' activate learners' epistemic agency, eventually leading to perceptual autonomy, which is the ability to identify cognitive weaknesses without AI assistance.

Challenges and Added Value

LMS such as Canvas and Blackboard have become nearly ubiquitous in higher education, with approximately 99% of colleges and universities in the United States employing some form of LMS (Dahlstrom et al., 2014, as cited in Kimberling & Akwafuo, 2023). Their built-in discussion boards represent a standard approach to facilitating asynchronous academic discourse, allowing instructors and students to post questions and responses at their own pace. This is an arrangement that, in theory, affords learners more time to read, reflect, and critically engage with course content before contributing (Kimberling & Akwafuo, 2023). Research confirms that these platforms support basic interaction and community-building in online courses, and they remain the dominant infrastructure for student-to-student and student-to-instructor communication (Han & Ellis, 2019, as cited in Kimberling & Akwafuo, 2023). However, critical analyses of LMS pedagogy reveal a structural tension: platforms like Canvas and Blackboard tend to reinforce a transmission model of education, in which content flows from instructor to student, rather than supporting the kind of open, dialogic, and problem-posing discourse that fosters deeper learning (Swerzenski, 2021). A campus-wide study of Canvas discussion activity further found that volume of posts alone does not predict student satisfaction or learning quality, underscoring that participation quantity and substantive intellectual engagement are fundamentally different outcomes (Rodriguez et al., 2021).

The persistent gap between the intended and actual outcomes of LMS discussion tools motivates the need for a new technological intervention. Despite their widespread adoption, conventional discussion boards frequently fall short of their pedagogical promise: students report that interactions feel repetitive and artificial, and research shows that while standard online discussion forums may improve some dimensions of critical reflection, they consistently underperform in supporting higher-order skills such as justification, critical assessment, and the integration of new ideas into prior knowledge (Al-Husban, 2020, as cited in ERIC, 2022; Mooney et al., 2014, as cited in Oregon State, 2022). Instructors must also invest significant effort to monitor, redirect, and scaffold discussions. This is a burden that scales poorly in large enrollment courses (Rodriguez et al., 2021). There are recent approaches to supporting knowledge synthesis including student-driven synthesis artifacts (Zhu et al., 2026) and LLM-based discourse modeling (Shui & Zhu, 2026). These studies have advanced the field but remain limited by the cognitive burden of manual creation and by post-hoc timing that cannot redirect active participation. DiscourseMap addresses these constraints through automated, real-time intervention within the existing LMS interface.

DiscourseMap addresses several challenges inherent in current graduate-level asynchronous discussions while adding significant value to the learning experience:

Table 1
Key Challenges & Value Propositions of DiscourseMap

Feature	Challenges Addressed	Added Value to Learning Experience	Supporting Literature
Epistemic Visibility	Standard LMS forums present posts as flat, chronological threads, making it difficult to see where ideas converge or diverge.	Transforms linear threads into a dynamic spatial map that renders the “epistemic landscape” visible in real time.	Wise & Cui (2018) Shui & Zhu (2026)
Thematic Clustering	Students often engage in “parallel monologues” or choose reply targets based on convenience rather than intellectual merit.	Uses AI-driven clustering to help learners identify thematic relationships and “logical flow” across the entire discussion.	Dillenbourg et al. (2009)
Cognitive Load Reduction	Text-heavy interfaces create a high cognitive burden for interpreting complex academic dialogues.	Employs multimodal visualization (network maps + text) to help students navigate argumentation pathways more efficiently.	Jang & Lee (2024)
Purposeful Interaction	Participation often becomes a “knowledge-telling” exercise to meet minimum post requirements.	Highlights “gaps” and “logical collision points,” providing context-sensitive prompts to trigger deeper inquiry and synthesis.	Scardamalia & Bereiter (2006)

Design Positioning

DiscourseMap is strategically positioned as a complementary and compensatory intervention that enhances existing educational infrastructure rather than replacing it. While established tools like Packback focus on inquiry-based questioning and Yellowdig emphasizes gamified community building, DiscourseMap distinguishes itself by focusing on the spatial synthesis of collective knowledge. As a complementary tool, it is implemented as a Canvas plugin, leveraging the fact that, as mentioned previously, 99% of higher education institutions already utilize an LMS (Hamadi et al., 2023). This allows it to augment existing discussion boards without requiring a full platform migration. It serves a compensatory function by addressing the "transmission model" of education often reinforced by standard LMS interfaces, which frequently fail to support higher-order skills like critical assessment and idea integration (Swerzenski, 2021). Furthermore, it substitutes linear navigation with a graph-driven approach, providing a level of "teaching presence" traditionally difficult to maintain in large-scale graduate seminars (Hullman et al., 2013). By acting as an "inter-objective artifact," DiscourseMap mediates between individual student thinking and the collective understanding of the group, providing a pedagogical scaffold that supports self-directed learning and metacognitive awareness (Owan et al., 2025).

Three core design decisions operationalize this positioning. First, the map is conceived as an analysis view rather than a replacement forum: students continue to read and post within the familiar Canvas thread interface, while the map provides navigational context and epistemic orientation. Replacing the posting environment would increase cognitive overhead, so the system overlays the existing interface. By preserving the Canvas shell, DiscourseMap acts as a layer of epistemic infrastructure that makes discourse structure visible without restructuring the discourse itself. This is consistent with Bodemer et al.'s (2018) principle that group awareness tools should make collective cognitive states legible to learners within, rather than apart from, the environment in which they work.

Second, gap clusters are designed to be actionable rather than indicative. In standard LMS environments, missing perspectives are structurally invisible in LMS discussions. There are no signals to learners that a theme is missing, underexplored, or uncontested. DiscourseMap materializes these absences as dashed-outline nodes on the spatial map, and clicking a gap node opens a reply flow specifically seeded for first contributions in that theme. This changes the gap from a passive analytical observation into an active invitation for knowledge-building participation (Scardamalia & Bereiter, 2006). The AI simulation showed that the tool successfully pointed out missing topics to all participants. However, behavioral uptake remained concentrated among higher-engagement personas, suggesting that gap actionability alone is insufficient for minimum compilers and that future iterations should investigate more direct scaffolding pathways toward gap engagement.

Third, epistemic prompts are designed to be context-sensitive. The prompts adapt to reply context: replying to a specific post activates cluster-aware prompts calibrated to that thematic region; replying to the main discussion activates map-wide prompts oriented toward global discourse structure; and replying to a gap activates gap-specific prompts that position the contribution as opening new epistemic territory. This scaffolding design operationalizes Ba et al.'s (2025) finding that AI-generated feedback is most effective when it is immediate, personalized, and calibrated to the learner's current epistemic position.

Taken together, these decisions reflect a consistent design philosophy: DiscourseMap intervenes at the moment of contribution rather than after it, making structural features of discourse visible without removing learner agency. The limitations of the current prototype (the absence of a live AI backend, the prototype-authored nature of cluster layouts and connections, and the deterministic character of synthetic replay) are intentional constraints appropriate to this stage of iterative design. They preserve the evaluative clarity needed to assess whether the core design mechanisms function as intended before introducing the additional variables of live AI inference. Future iterations will address these constraints as the tool moves toward deployment with human participants under IRB-approved conditions.

Synthetic Formative Evaluation

Theoretical Grounding for Persona and Scenario Design

The three epistemic personas — Minimum Complier, Surface Engager, and Curious Builder — were derived from Van Aalst's (2009) empirical distinction between knowledge-sharing, knowledge-construction, and knowledge-creation discourse. Van Aalst demonstrated that participation quality in online discussions is determined not by volume but by the nature of contributions: whether learners restate existing understanding, extend and interpret ideas, or generate genuinely new conceptual connections. Each persona represents a distinct discourse mode (sharing, construction, creation), allowing the evaluation to assess whether DiscourseMap's design mechanisms produced differential behavioral outcomes across the full participation spectrum.

The three scenario conditions (Surface, Deep, and Polarized) were grounded in Scardamalia and Bereiter's (2006) account of the structural conditions that enable or inhibit knowledge building. Scardamalia and Bereiter distinguish between environments that support knowledge-telling, in which participants express existing understanding without advancing collective inquiry, and those that support knowledge-building, in which discourse is organized around the progressive improvement of ideas. The Surface scenario was designed to approximate the knowledge-telling baseline that LMS discussion forums typically produce (Wise et al., 2014); the Deep scenario approximated an idealized knowledge-building environment; and the Polarized scenario introduced the additional structural condition of premature consensus within competing camps, a pattern that Stahl and Hakkarainen (2021) identify as a characteristic failure mode of computer-supported collaborative learning environments in which epistemic structure is not made visible to participants. Together, personas and scenarios test the system across varying levels of engagement and polarization.

Overview

Three discussion scenarios were constructed to represent distinct epistemic states: a surface-level discussion characterized by consensus convergence and minimal cross-cluster engagement (Surface), a knowledge-building discussion featuring substantive cross-cluster bridging (Deep), and a polarized discussion in which two dominant ideological clusters formed with no bridging between them (Polarized). Each scenario included ten posts from five synthetic participants.

To simulate ecologically valid participation patterns, each student was assigned a fixed epistemic persona drawn from the Knowledge Building literature: Jennifer was assigned as a Minimum Complier, Stella and Liz as Surface Engagers, and Addison and Matty as Curious Builders. Students participated in a staggered turn order — Jennifer and Addison posting first, followed by Stella, Matty, and Liz — so that each subsequent participant could respond to an evolving discussion state. Session events were logged automatically, capturing cluster navigation, prompt engagement, reply content, and cross-cluster behavior. All three sessions were exported as structured JSON files for analysis.

Behavioral Findings

Table 2 presents map exploration and scaffolding engagement data across all 15 student-turn combinations (5 students × 3 scenarios).

Table 2
Map Exploration and Scaffolding Engagement by Persona and Scenario

Student	Persona	Scenario	Clusters Visited	Posts Read	Gap Engaged	Prompts Time (ms)	Prompt_used	Cross Cluster	Reply Length (chars)
Jennifer	Min. Complier	Surface	1	0	X	303	X	X	244
Jennifer	Min. Complier	Deep	1	1	X	302	X	X	273
Jennifer	Min. Complier	Polarized	1	1	X	303	X	X	293
Stella	Surface Engager	Surface	1	2	X	603	X	X	524
Stella	Surface Engager	Deep	1	3	X	604	O	X	636
Stella	Surface Engager	Polarized	1	3	X	602	X	X	687
Liz	Surface Engager	Surface	1	3	X	603	X	X	558
Liz	Surface Engager	Deep	1	3	X	603	O	X	875
Liz	Surface Engager	Polarized	1	2	X	602	X	X	816
Addison	Curious Builder	Surface	5	5	O (instructional-design)	904	O	O	1,749
Addison	Curious Builder	Deep	4	6	O (student-agency)	—	O	O	1,694
Addison	Curious Builder	Polarized	4	4	O (hybrid-models)	—	O	O	1,772
Matty	Curious Builder	Surface	5	2	O (equity-access)	—	O	O	1,630

Matty	Curious Builder	Deep	5	4	O (student -agency)	—	X	O	1,761
Matty	Curious Builder	Polarized	3	6	O (student -perspective)	—	X	O	1,539

Key Findings

Finding 1: Prompt Dismissal Time Reliably Indexes Scaffolding Engagement. Across all three scenarios, Jennifer dismissed epistemic prompts within 302–303ms — a duration insufficient to read a single prompt card. Stella and Liz dismissed prompts within 602–604ms, consistent with brief exposure but without active selection. In contrast, Addison and Matty triggered prompt_inserted events, indicating they selected and inserted at least one prompt starter before composing their replies. Notably, promptStarterUsed remained false across all sessions, indicating that scaffolding influenced the direction of contribution without determining its content — a pattern consistent with productive rather than prescriptive scaffolding effects.

Finding 2: Gap Cluster Engagement Was Exclusive to Curious Builder Personas. In every scenario, Addison and Matty navigated to at least one gap cluster (isGap: true) and posted there, generating isCrossCluster: true events with multi-cluster attribution spanning two to three thematic regions. Jennifer, Stella, and Liz posted exclusively within a single pre-existing cluster in all sessions (isCrossCluster: false). This finding provides evidence that DiscourseMap's gap visibility mechanism activated differential behavior by epistemic orientation — the gap indicators were perceptually available to all participants, but only Curious Builder personas translated that visibility into navigational action.

Finding 3: The Deep Scenario Produced the Most Distributed Epistemic Structure. Across all three scenarios, the Deep session generated 23 total posts across 5 fully populated clusters, with 0 remaining gaps and 0 consensus risks at session close. The Polarized session produced the most persistent structural tension: despite Addison and Matty bridging into Hybrid Approaches and Student Perspective respectively, Policy & Governance remained unpopulated and both dominant clusters retained consensus warnings. The Surface session eliminated both initial gap clusters but retained a consensus warning in Cognitive Offloading, reflecting the difficulty of dislodging an entrenched dominant frame even with gap-directed contributions.

Finding 4: The extra_themes_set Event Reveals DiscourseMap's Bridging Mechanism in Action. When Addison and Matty posted into gap clusters, both consistently triggered extra_themes_set events attributing their posts to two or three clusters simultaneously. In the Surface scenario, Addison's instructional-design post was attributed to [instructional-design, cognitive-offloading, metacognition-ai]; in the Polarized scenario, Addison's hybrid-models post was attributed to [hybrid-models, human-relation-irreplaceable, ai-tutor-benefits]. This cross-cluster attribution is the structural mechanism by which DiscourseMap transforms isolated gap clusters into connected regions — precisely the knowledge-building behavior the tool is designed to support (Scardamalia & Bereiter, 2006).

Implications for Design

These findings offer three design-relevant conclusions. First, the gap indicator mechanism successfully surfaced underexplored areas to all participants, but behavioral uptake was contingent on epistemic orientation — suggesting that gap visibility alone is insufficient and that prompt scaffolding may need to more actively invite lower-engagement personas toward gap clusters. Second, the prompt engagement pattern (visible but not prescriptive) is consistent with effective scaffolding design: starters were used as directional cues rather than templates, preserving epistemic agency. The dissociation between prompt_inserted and promptStarterUsed is particularly informative — it suggests that the prompts encourage students to think for themselves rather than just writing for them, which aligns with productive scaffolding theory. Third, the persistence of consensus warnings in the Polarized scenario even after bridging contributions suggests that DiscourseMap's consensus warning threshold may need recalibration — a finding that informs the design refinement cycle prior to formal user testing.

User Testing

Participants and Context

A preliminary user test was conducted with six Columbia University graduate students (P1–P6), all of whom had prior experience with asynchronous LMS discussion boards as a course requirement. Participants were recruited informally and were not compensated. Because this evaluation preceded formal IRB review, it was designed as an exploratory usability study rather than a controlled experiment; the goal was to assess whether the core design mechanisms of DiscourseMap functioned as intended under authentic graduate-student usage conditions. The discussion topic was Week 12 — AI Narrative Intelligence: Creative Agency and Data Storytelling, grounded in two assigned readings on AI and data visualization (Ciuccarelli, 2023) and AI narrative techniques (Flash, 2024). The live discussion produced 14 posts across 5 active thematic clusters, with 1 gap cluster identified and 3 clusters flagged as needing counter-perspectives.

Anonymization and Study Setup

To minimize social desirability effects and encourage naturalistic participation, each participant was assigned a unique study code at login. Within the shared discussion thread, all contributions were displayed under anonymized identifiers (TC-XXX) rather than real names, so that participants could engage freely without social pressure or self-presentation concerns. This design decision reflects the principle that epistemic risk-taking — posting tentative ideas, challenging peers, contributing to underexplored gaps — is suppressed when identity is salient (Kipling et al., 2023). For the same reason, session logs and survey responses were analyzed at the group level only; individual-level matching was not conducted to preserve participant anonymity.

Technical Infrastructure

The study session ran on the Firebase-backed version of DiscourseMap, enabling real-time shared state across all participants: posts, replies, and theme configurations were synchronized live, so that the cluster map and discussion health indicators updated dynamically as the discussion evolved. This condition approximated an authentic collaborative discussion environment rather than an isolated single-user prototype test.

Data Collection

Three complementary data streams were collected. First, meticulous behavioral logs were recorded locally for each participant using the prototype's built-in event logging system. Participants initiated logging via a Start Log control; from that point, the system captured all interactions — cluster selections, hover dwells, prompt openings, prompt dismissals, reply composition, and submission — as timestamped JSON events. Second, a post-session survey was administered immediately following the session, covering prior LMS discussion experience, feature usage, perceived utility across five dimensions, and open-ended reflections. Third, real-time observation of the shared Firebase thread allowed the research team to monitor discussion evolution, cluster population, and gap engagement across participants during the session.

User Testing Results

Behavioral Logs. All participants successfully navigated the interaction logging process without technical difficulty. Table 4 presents session log metrics for all six participants.

Table 4
Logs for six participants

P	Clusters Visited	Keywords Hovered	Gap Cluster Viewed	Conne ctions	Prompts Viewed	Prompt Inserted	Replies Submitted	Avg. Reply Length	Prompt Starter Used	Gap Posted
P1	1	18	4	1	3	0	1	1,226	X	X
P2	2	49	1	1	15	9	2	1,014	O	X
P3	0	0	0	0	9	0	3	630	X	X
P4	2	32	2	2	12	2	4	483	O	X
P5	5	52	3	3	18	9	2	1,122	X	X
P6	3	17	1	0	6	2	2	464	O	O
Mean	2.2	28.0	1.8	1.2	10.5	3.7	2.3	823	3/6	1/6

Map exploration varied substantially: cluster navigation ranged from 0 (P3) to 5 (P5), with a group mean of 2.2. P3 registered no map interaction whatsoever — yet submitted 3 replies — suggesting DiscourseMap was used as a writing scaffold rather than an epistemic map. P5 generated the richest trace (312 total events), selected all 5 clusters, hovered 52 keywords, and produced the second-highest mean reply length (1,122 chars). P1 navigated only 1 cluster but hovered 18 keywords, suggesting that keyword chips served as an alternative entry point into the discourse landscape for some users.

Five of six participants (83%) viewed at least one gap cluster; only P6 submitted a reply into a gap (policy-privacy-power), using a prompt starter to initiate the contribution. This high-visibility, low-uptake pattern — gaps are noticed but rarely acted upon — was consistent across all participants except P6. Four participants used prompt_inserted at least once; three submitted at least one reply with promptStarterUsed recorded as true. In all three cases the starter functioned as a launching phrase extended with original content rather than a template, consistent with the expectation that prompts should influence direction without determining content. One exception was P4's 95-character reply that used a starter without meaningful elaboration. No participant submitted a cross-cluster reply (isCrossCluster: false for all 6).

Post-Session Survey. Participants reported high levels of satisfaction (scoring 4 or 5 out of 5), confirming that DiscourseMap effectively helped them visualize discussion structures and identify meaningful opportunities for contribution. Furthermore, the survey results indicate that the tool successfully encouraged users to recognize thematic connections and engage in deeper, more thoughtful discourse. In addition, the survey provided qualitative insights into how students interacted with DiscourseMap's core design mechanisms. First, the user testing results revealed that while DiscourseMap significantly enhanced structural visibility and reduced cognitive load through its "visual map" and "AI-generated guides", it also presented specific challenges. Participants appreciated the system's ability to clarify discussion landscapes but noted that early exposure to it could lead to "idea fixation," potentially constraining independent thinking. To address these limitations, the participants suggested key areas for optimization, such as implementing a staged disclosure of the map (revealing it only after the initial post) and refining the Theme View to extract cross-post granular data. Second, two participants independently reported not understanding the red exclamation mark (consensus warning) badges, suggesting a discoverability gap that affects one of DiscourseMap's core epistemic signals.

Table 5.

Post-Session Survey Likert Results (n=6, scale 1–5)

Survey Item (1–5 scale)	Min	Max	Mean	n=6
Q10. DiscourseMap helped me understand the overall structure of the discussion.	4	5	4.17	6
Q11. DiscourseMap helped me identify where I could contribute meaningfully.	3	5	4.33	6
Q12. DiscourseMap helped me see connections between different ideas.	4	5	4.83	6
Q13. DiscourseMap encouraged me to go beyond simple agreement.	4	5	4.50	6
Q14. The tool helped reduce the cognitive load of reading long threads.	3	5	4.17	6

Iterative Design Implications. The gap between visibility and action converged across both conditions: gaps were viewed but rarely acted upon regardless of whether participants were synthetic or real. Cross-cluster posting, which occurred in all Curious Builder synthetic sessions, did not occur in any live session, suggesting that the current design supports epistemic orientation but not yet cross-cluster synthesis. The prompt starter pattern was also consistent: starters functioned as triggers rather than templates in both conditions, the most direct confirmation that DiscourseMap's scaffolding mechanism operates as designed. Therefore, three issues warrant attention in future iterations. First, gap indicators require stronger behavioral affordances — five participants viewed gaps without acting on them, suggesting that the visual signal alone is insufficient to redirect participation. Second, consensus warning badges require improved legibility — two participants did not understand their meaning, undermining one of the tool's core epistemic signals. Third, cross-cluster synthesis requires explicit invitation — map visibility alone did not produce bridging behavior in any live session, pointing to the need for prompts that are specifically triggered when a participant has navigated multiple clusters before opening a reply.

Product Overview

DiscourseMap (https://haetalkim.github.io/discourse-map/)

DiscourseMap uses AI-driven clustering to transform the flat chronological thread of an LMS discussion into a spatial representation of its thematic landscape, rendering visible the distribution, density, and relationships among ideas in real time. Unlike post-hoc discourse tools—such as learning analytics dashboards (Wise & Cui, 2018) or LLM-generated knowledge synthesis graphs (Shui & Zhu, 2026)—this visualization updates as new contributions are added, functioning as what Stahl and Hakkarainen (2021) describe as an inter-objective artifact that mediates between individual and collective understanding during active participation.

The system incorporates three scaffolding mechanisms as mentioned earlier. First, as shown in Figures 1 and 2, *Discourse visualization* makes the collective thinking visible by mapping thematic clusters and their connections. Second, *Gap indicators*—integrated within the same visual mapping— extend this mapping to highlight underexplored areas, missing counter-perspectives, and clusters lacking critical engagement, drawing on evidence that group awareness of cognition improves coordination (Bodemer et al., 2018), while addressing the limits of student-led identification of overlooked ideas (Zhu et al., 2026). Third, *Epistemic prompts* (Figure 3) then translate insights into actionable guidance at the level of collective discourse rather than the individual posts (cf. Ba et al., 2025), responding to the finding that synthesis instructions alone do not ensure epistemic depth (Zhu et al., 2026). Implemented as a Canvas LMS plugin, the tool embeds these mechanisms directly into the existing discussion interface, minimizing adoption barriers (Dillenbourg et al., 2009).

Figure 1
Overview of the DiscourseMap Interface and Core Visual Components

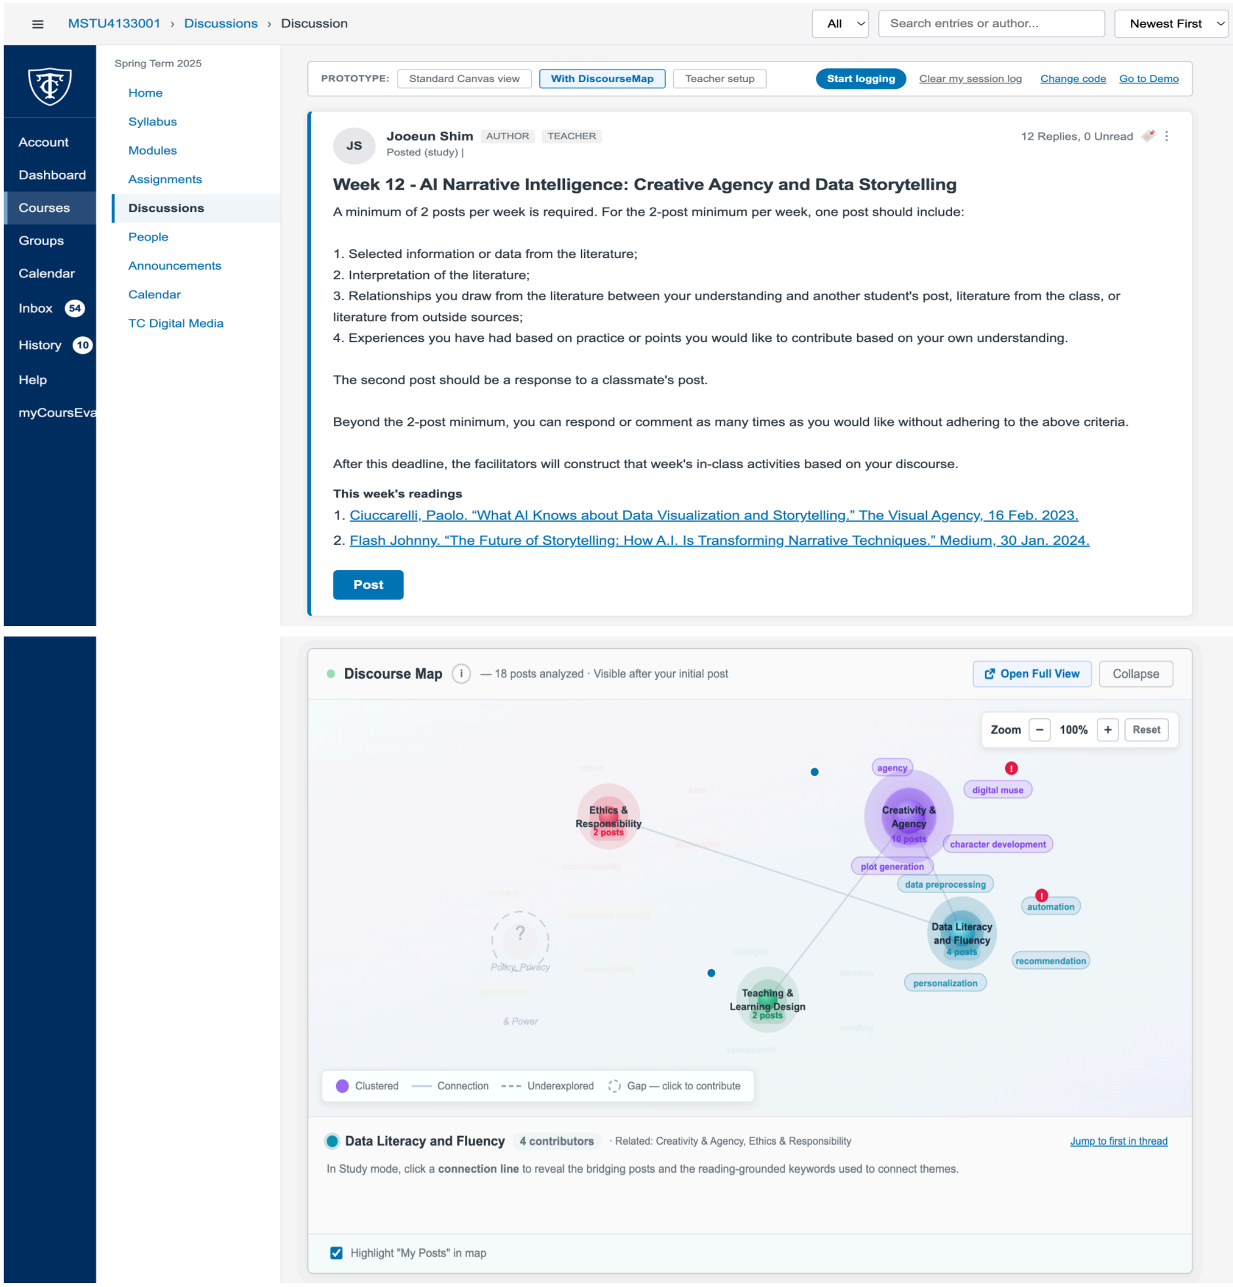

MSTU4133001 › Discussions › Discussion
All
Search entries or author...
Newest First
Account
Dashboard
Courses
Groups
Calendar
Inbox 54
History 10
Help
myCoursEva
Spring Term 2025
Home
Syllabus
Modules
Assignments
Discussions
People
Announcements
Calendar
TC Digital Media
PROTOTYPE: Standard Canvas view With DiscourseMap Teacher setup
Start logging Clear my session log Change code Go to Demo
JS
Jooeun Shim AUTHOR TEACHER
Posted (study) |
12 Replies, 0 Unread
Week 12 - AI Narrative Intelligence: Creative Agency and Data Storytelling
A minimum of 2 posts per week is required. For the 2-post minimum per week, one post should include:
1. Selected information or data from the literature;
2. Interpretation of the literature;
3. Relationships you draw from the literature between your understanding and another student's post, literature from the class, or literature from outside sources;
4. Experiences you have had based on practice or points you would like to contribute based on your own understanding.
The second post should be a response to a classmate's post.
Beyond the 2-post minimum, you can respond or comment as many times as you would like without adhering to the above criteria.
After this deadline, the facilitators will construct that week's in-class activities based on your discourse.
This week's readings
1. Ciuccarelli, Paolo. "What AI Knows about Data Visualization and Storytelling," The Visual Agency, 16 Feb. 2023.
2. Flash Johnny. "The Future of Storytelling: How A.I. Is Transforming Narrative Techniques." Medium, 30 Jan. 2024.
Post
Discourse Map — 18 posts analyzed · Visible after your initial post
Open Full View
Collapse
Zoom 100% Reset
agency
digital muse
Ethics & Responsibility 2 posts
Creativity & Agency 10 posts
character development
plot generation
data preprocessing
automation
Data Literacy and Fluency 4 posts
recommendation
personalization
Policy, Privacy & Power
Teaching & Learning Design 2 posts
Clustered Connection Underexplored Gap — click to contribute
Data Literacy and Fluency 4 contributors · Related: Creativity & Agency, Ethics & Responsibility
Jump to first in thread
In Study mode, click a connection line to reveal the bridging posts and the reading-grounded keywords used to connect themes.
Highlight "My Posts" in map

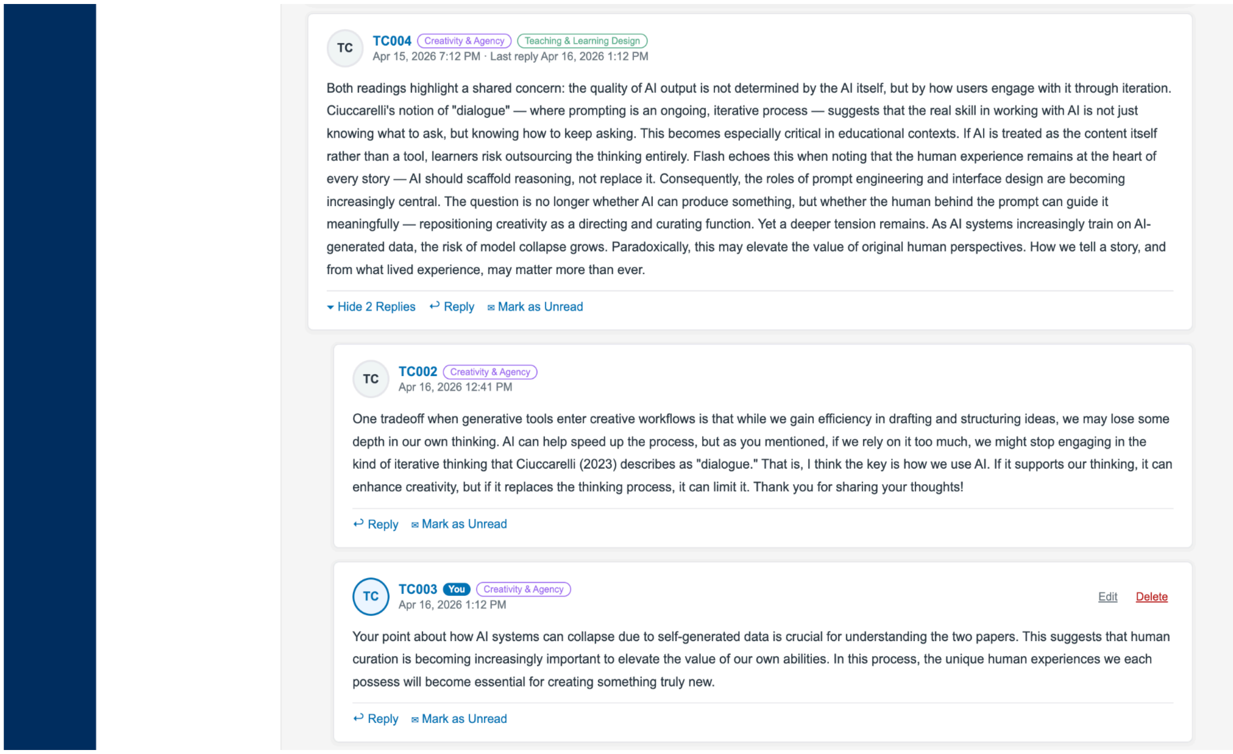

Figure 2

Full View Mode of the Discourse Map

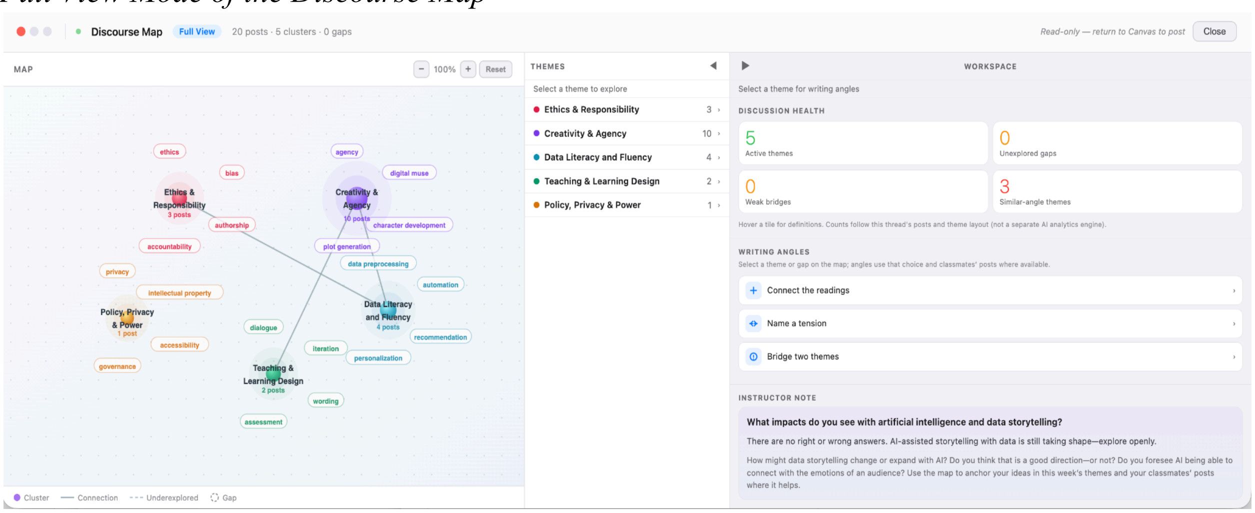

Figure 3

AI Writing Prompt for Posting and Replying

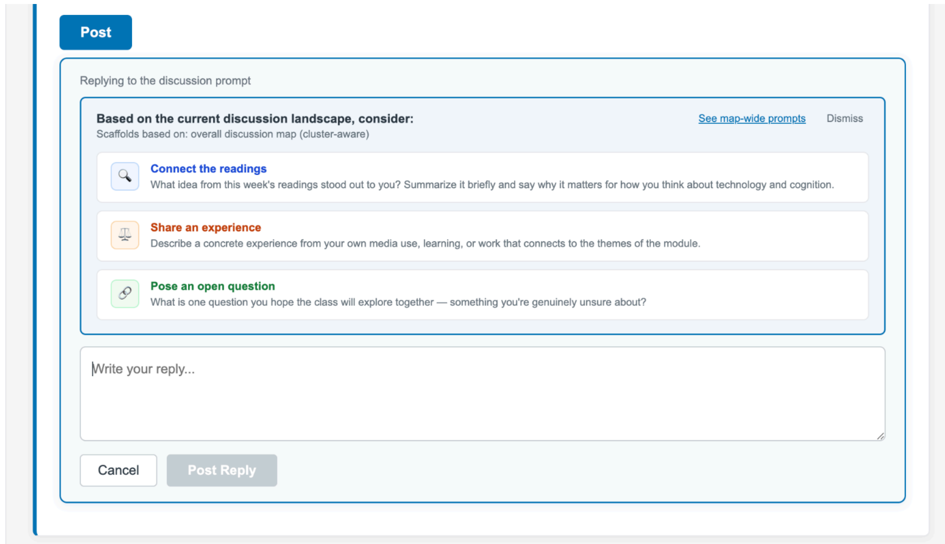
Post
Replying to the discussion prompt
Based on the current discussion landscape, consider:
Scaffolds based on: overall discussion map (cluster-aware)
See map-wide prompts
Dismiss
Connect the readings
What idea from this week's readings stood out to you? Summarize it briefly and say why it matters for how you think about technology and cognition.
Share an experience
Describe a concrete experience from your own media use, learning, or work that connects to the themes of the module.
Pose an open question
What is one question you hope the class will explore together — something you're genuinely unsure about?
Write your reply...
Cancel
Post Reply

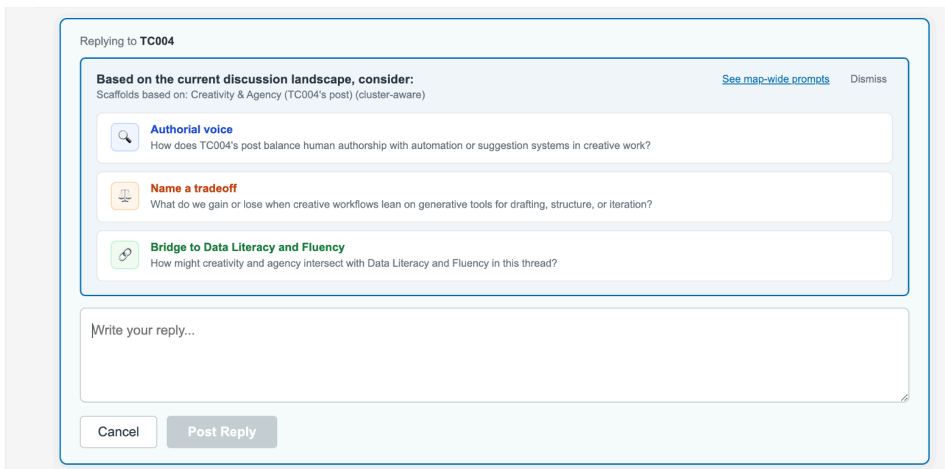
Replying to TC004
Based on the current discussion landscape, consider:
Scaffolds based on: Creativity & Agency (TC004's post) (cluster-aware)
See map-wide prompts
Dismiss
Authorial voice
How does TC004's post balance human authorship with automation or suggestion systems in creative work?
Name a tradeoff
What do we gain or lose when creative workflows lean on generative tools for drafting, structure, or iteration?
Bridge to Data Literacy and Fluency
How might creativity and agency intersect with Data Literacy and Fluency in this thread?
Write your reply...
Cancel
Post Reply

Reference

Abdulbaki, K., Suhaimi, M., Alsaqqaf, A., & Jawad, W. (2018). The use of the discussion method at university: Enhancement of teaching and learning. *International Journal of Higher Education, 7*(6), 118–128.

Asterhan, C. S. C. (2015). Introducing online dialogues in collocated classrooms: If, why, and how. In L. B. Resnick, C. S. C. Asterhan, & S. N. Clarke (Eds.), *Socializing intelligence through academic talk and dialogue* (pp. 205–218). American Educational Research Association. https://doi.org/10.3102/978-0-935302-43-1_16

Ba, S., Zhan, Y., Huang, L., & Lu, G. (2025). Investigating the impact of ChatGPT-assisted feedback on the dynamics and outcomes of online inquiry-based discussion. *British Journal of Educational Technology, 56*(6), 1710–1734. https://doi.org/10.1111/bjet.13605

Barzilai, S., & Chinn, C. A. (2018). On the goals of epistemic education: Promoting apt epistemic performance. *Journal of the Learning Sciences, 27*(3), 353–389. https://doi.org/10.1080/10508406.2017.1392968

Bodemer, D., Janssen, J., & Schnaubert, L. (2018). Group awareness tools for computer-supported collaborative learning. In F. Fischer, C. E. Hmelo-Silver, S. R. Goldman, & P. Reimann (Eds.), *International handbook of the learning sciences* (pp. 351–358). Routledge. https://doi.org/10.4324/9781315617572-34

Cal-Varela, M., & Fernández-Polo, F. J. (2022). Referring to other participants in asynchronous online discussions: Citation patterns in a higher education context. *Psychology of Language and Communication, 26*(1), 353–374.

Dillenbourg, P., Järvelä, S., & Fischer, F. (2009). The evolution of research on computer-supported collaborative learning. In N. Balacheff, S. Ludvigsen, T. de Jong, A. Lazonder, & S. Barnes (Eds.), *Technology-enhanced learning: Principles and products* (pp. 3–19). Springer. https://doi.org/10.1007/978-1-4020-9827-7_1

Dovhaniuk, E., Piwowar, A., & Oliynyk, N. (2025). Multimodal and cognitive approaches to academic discourse in AI-supported learning. *Cognition, Communication, Discourse, (30)*, 24–39. https://doi.org/10.26565/2218-2926-2025-30-02

Gasell, C., Lowenthal, P. R., Uribe-Flórez, L. J., & Ching, Y. H. (2022). Interaction in asynchronous discussion boards: A campus-wide analysis to better understand regular and substantive interaction. *Education and Information Technologies, 27*(3), 3421–3445.

Harati, H., Ono, E., Ingram, J. K., Nasr, N. R., Yen, C. J., & Tu, C. H. (2023). Exploring the impact of discussion interfaces on diversified interconnectivity in online learning communities. *Journal of Educational Technology Development and Exchange (JETDE), 16*(2), 40–65.

Jang, M., & Lee, H. (2024). Using topic modeling-based network visualization and generative AI in online discussions: How learners' perception of usability affects their reflection on feedback. *Educational Technology International, 25*(1), 1–25. https://doi.org/10.23095/ETI.2024.25.1.001

Kauer, T., Akbaba, D., Dörk, M., & Bach, B. (2024). Discursive patinas: Anchoring discussions in data visualizations. *IEEE Transactions on Visualization and Computer Graphics, 31*, 1246–1256. https://doi.org/10.1109/TVCG.2024.10679900

Kipling, R. P., Stiles, W. A., de Andrade-Lima, M., MacKintosh, N., Roberts, M. W., Williams, C. L., & Watson-Jones, S. J. (2023). Interaction in online postgraduate learning: What makes a good forum? *Distance Education, 44*(1), 162–189.

Long, Z., Luo, D., Kiu, K., Gao, H., Qu, J., & Hu, X. (2020). Raising academic performance in socio-cognitive conflict learning through gamification. In *International Conference on Artificial Intelligence in Education* (pp. 180–184). Springer. https://doi.org/10.1007/978-3-030-52240-7_33

O'Brien, J. F. (1939). A definition and classification of the forms of discussion. *Quarterly Journal of Speech, 25*(2), 236–243.

O'Connell, T., & Dyment, J. (2006). Reflections on using journals in higher education: A focus group discussion with faculty. *Assessment & Evaluation in Higher Education, 31*, 671–691. https://doi.org/10.1080/02602930600760884

Parker, W. C., & Hess, D. (2001). Teaching with and for discussion. *Teaching and Teacher Education, 17*(3), 273–289.

Reddy, Y., & Andrade, H. (2010). A review of rubric use in higher education. *Assessment & Evaluation in Higher Education, 35*, 435–448. https://doi.org/10.1080/02602930902862859

Reed, I., & Alexander, J. (2009). Social science as reading and performance: A cultural-sociological understanding of epistemology. *European Journal of Social Theory, 12*(1), 21–41. https://doi.org/10.1177/1368431008099648

Rusen, M. (2025). Discourse in online learning environments: A systematic review of contemporary research literature. *Open Access Journal of Education & Language Studies, 3*(1), Article 555602. https://doi.org/10.19080/OAJELS.2025.03.555602

Scardamalia, M., & Bereiter, C. (2006). Knowledge building: Theory, pedagogy, and technology. In K. Sawyer (Ed.), *The Cambridge handbook of the learning sciences* (pp. 97–118). Cambridge University Press. https://doi.org/10.1017/CBO9780511816833.008

Shea, P., Li, C. S., & Pickett, A. (2006). A study of teaching presence and student sense of learning community in fully online and blended graduate courses. *The International Review of Research in Open and Distributed Learning, 7*(2). https://doi.org/10.19173/irrodl.v7i2.290

Shui, B., & Zhu, X. (2026). Knowledge synthesis graph: An LLM-based approach for modeling student collaborative discourse. *arXiv*. https://doi.org/10.48550/arXiv.2602.06194

Stahl, G., & Hakkarainen, K. (2021). Theories of CSCL. In U. Cress, C. Rosé, A. F. Wise, & J. Oshima (Eds.), *International handbook of computer-supported collaborative learning* (pp. 23–43). Springer. https://doi.org/10.1007/978-3-030-65291-3_2

Van Aalst, J. (2009). Distinguishing knowledge-sharing, knowledge-construction, and knowledge-creation discourses. *International Journal of Computer-Supported Collaborative Learning, 4*(3), 259–287.

Wise, A. F., Zhao, Y., & Hausknecht, S. N. (2014). Learning analytics for online discussions: Embedded and extracted approaches. *Journal of Learning Analytics, 1*(2), 48–71. https://doi.org/10.18608/jla.2014.12.4

Ying, J. (2020). The importance of the discussion method in the undergraduate business classroom. *Humanistic Management Journal, 5*(2), 251–278.

Yoshida, M., Duangchinda, V., & Ruangrit, N. (2026). AI-supported learning in online discussion forums: A scoping review. *Electronic Journal of e-Learning, 24*(1), 19–35.

Zhu, X., Chen, B., & DeLiema, D. (2026). Advancing collaborative discourse through knowledge synthesis. *International Journal of Computer-Supported Collaborative Learning*. Advance online publication. https://doi.org/10.1007/s11412-025-09463-6