ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

Learning Companion for students' critical reasoning and academic writing

Chuan You Tan,
Nanyang Technological University, Singapore

Owen Noel Newton Fernando
Nanyang Technological University, Singapore

Jun Hao Mark Tan
National Institute of Education, Singapore

Chin Ann Ong
Nanyang Technological University, Singapore

Timothy Robert Merritt
Aalborg University, Denmark

Victor Lim Fei
National Institute of Education, Singapore

The advent of Generative Artificial Intelligence (GenAI) transformed technology's role in the education sector by introducing highly adaptive, interactive tools. Standard implementations are often inadequately designed to effectively aid students. In some cases, these systems hinder genuine learning by inducing cognitive offloading, offering immediate solutions that bypass the cognitive processes necessary for deep conceptual understanding. This concern highlights the need for pedagogically grounded systems that foster genuine understanding and skills development. To address this gap, this paper presents an experimental implementation evaluating the feasibility of embedding formal educational frameworks directly into AI system architecture. We introduce Learning Companion, a prototype engineered to mitigate passive cognitive offloading by translating educational theory into programmatic workflows. Rather than replacing educators or automating learning, the Learning Companion is designed to strengthen meaningful learning interactions by positioning AI as a collaborative partner that supports learner agency through guided dialogue and adaptive instructional scaffolding. Aimed at enhancing critical reasoning and academic writing, the system operationalises Vygotsky's Zone of Proximal Development (ZPD) through algorithmic Socratic questioning to provide structured instructional scaffolding. This human-centred approach demonstrates how AI can complement teaching and foster deeper learner engagement while preserving meaningful connections between students, educators, and intelligent learning technologies.

Keywords: Artificial Intelligence, critical reasoning, human-centred learning, Socratic learning, zone of proximal development

Background

Digital technologies have fundamentally altered higher education by expanding interaction spaces and resource accessibility (Dhameria et al., 2025; Yang et al., 2024). Recently, the rapid innovation of Generative AI (GenAI) has transformed how we view technology's role in the education sector by introducing highly

adaptive, interactive tools (Crompton & Burke, 2023; Kasneci et al., 2023). However, standard implementations of these tools are often inadequately designed to effectively aid students. In many cases, consumer-grade GenAI configurations hinder genuine learning by inducing cognitive offloading, offering immediate solutions that bypass the cognitive means necessary for deep conceptual understanding and critical analysis (Abubakar et al., 2025; Gerlich, 2025; Jose et al., 2025). This problematic reliance threatens independent problem-solving and the development of core academic competencies expected of tertiary students, where critical reasoning serves as the pillar for academic writing success (Chukwuere, 2024).

This structural deficiency is deeply rooted in a broader historical misalignment within educational conversational agent development. A systematic review of 71 empirical studies spanning 2019 to 2023 revealed a profound "theory gap": 82% of existing educational chatbots were deployed without anchoring their design in a formal educational or learning theory, risking complete misalignment with actual human cognitive processes (Debets et al., 2025). Furthermore, these implementations have been structurally limited, focusing mostly on teaching-oriented outcomes (65%) or administrative logistics (31%), heavily concentrated in STEM fields (49%).

To address these theoretical and structural gaps, contemporary educational frameworks emphasize the need for "social generative AI" (Sharples, 2023) and customized dialogic partners capable of actively prompting student reasoning and multi-perspective argumentation (Tang & Putra, 2025). Rather than prioritizing immediate answer generation, educational technology must pivot toward architectures that translate educational theory into programmatic workflows. The web-based Learning Companion described in this paper mitigates passive cognitive offloading by operationalising Vygotsky's Zone of Proximal Development (ZPD) through Socratic questioning and structured instructional scaffolding. Implemented via a robust technical stack—featuring a React frontend served by NGINX, a Python Fast API backend, MongoDB, and the Dify orchestration framework—this system validates the feasibility of embedding formal educational frameworks directly into AI system architecture, thereby pivoting current GenAI practices to support human-centred learning design (Peláez et al., 2025).

Pedagogical framework

To bridge the empirical "theory gap" identified in educational technology literature, the design of the Learning Companion is explicitly grounded in two instructional models: Vygotsky's Zone of Proximal Development (ZPD) and Socratic Questioning. Rather than treating these concepts as abstract guidelines, they are operationalised directly into the system's programmatic interaction constraints.

Zone of proximal development (ZPD)

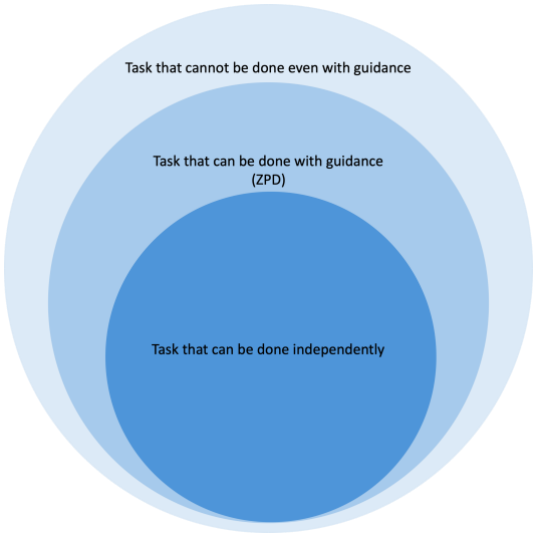


Figure 1. Illustration of Vygotsky's Zone of Proximal Development (ZPD)

The ZPD defines the cognitive space between what a learner can achieve independently and what they can accomplish with structured guidance (Vygotsky, 1978). Illustrated in Figure 1, tasks are separated into three classes where the inner circle depicts the comfort zone of the learners, the middle circle depicts the tasks that learner can achieve with some guidance and finally the outer circle represents tasks that are impossible to achieve even with guidance. The classic ZPD limits the More Knowledgeable Other (MKO) to human educators or peers. However, modern adaptation with technologies argues that technological artifacts can serve as the replacement. In our case, GenAI was considered as the MKO (Abtahi et al., 2017). Within this platform, scaffolding is dynamically calibrated to match the student's immediate level of reasoning playing the MKO role. Instead of taking over writing tasks or offering pre-formulated text generation shortcuts, the system operates exclusively within this middle range, providing micro-hints and analytical checkpoints that push the learner to resolve conceptual ambiguities independently.

Algorithmic Socratic questioning

To actively counter passive cognitive offloading, the system utilises Socratic dialogue as its primary conversational mechanism. When a student submits a text block or queries an uploaded document, the system does not only output answers. Instead, it executes an interrogation sequence across three instructional dimensions:

- Clarification Prompts: Compelling the student to define underlying assumptions or ambiguous terms within their text.
- Evidence Verification: Prompting the student to evaluate whether their cited premises are sufficient to defend their core thesis.
- Adversarial Counter Arguments: Intentionally presenting alternative interpretations or exposing logical fallacies to test the robustness of the student's stance.

To structurally enforce active learning, a mandatory system prompt constraint dictates that every single agent response must conclude with an open-ended, reflective question or task. This creates a continuous feedback loop that shifts the cognitive burden back onto the student, ensuring that the technology supports human-centred learning design (Favero et al., 2024).

Methodology

Moving from theory to execution, the system architecture maps these pedagogical rules into a deterministic, state-driven multi-agent workflow managed via the Dify orchestration framework (LangGenius, 2024).

Multi-Layer topology

The application leverages a modular three-layer architecture to cleanly decouple presentation mechanics, business logic, and database state management as shown in Figure 2.

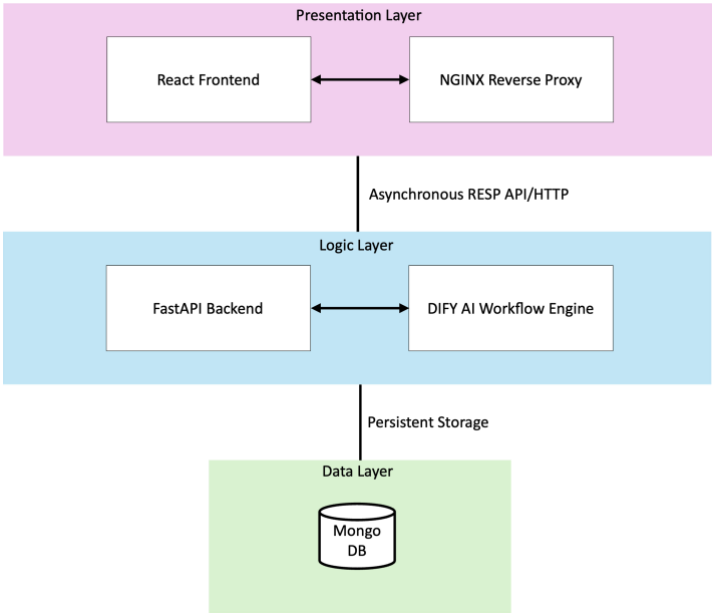


Figure 2. Architectural design of Learning Companion

Presentation Layer: A component based React user interface is served via a NGINX reverse proxy to maximise system delivery stability. The layout features a three-segment workspace: communication state history, a main processing workspace (incorporating an interactive PDF annotation canvas and a writing draft space), and a dedicated Socratic chat terminal.

Logic Layer: A Python Fast API backend serves as the core secure engine, orchestrating application authentication, pre-processing, and file routing operations. It passes structured data blocks to Dify via API protocols, establishing an isolated AI execution container (Ramirez, 2020).

Data Layer: A NoSQL MongoDB instance provides non-relational database storage for persistent user indexing and flexible schema management (Chodorow, 2013).

User interface and features

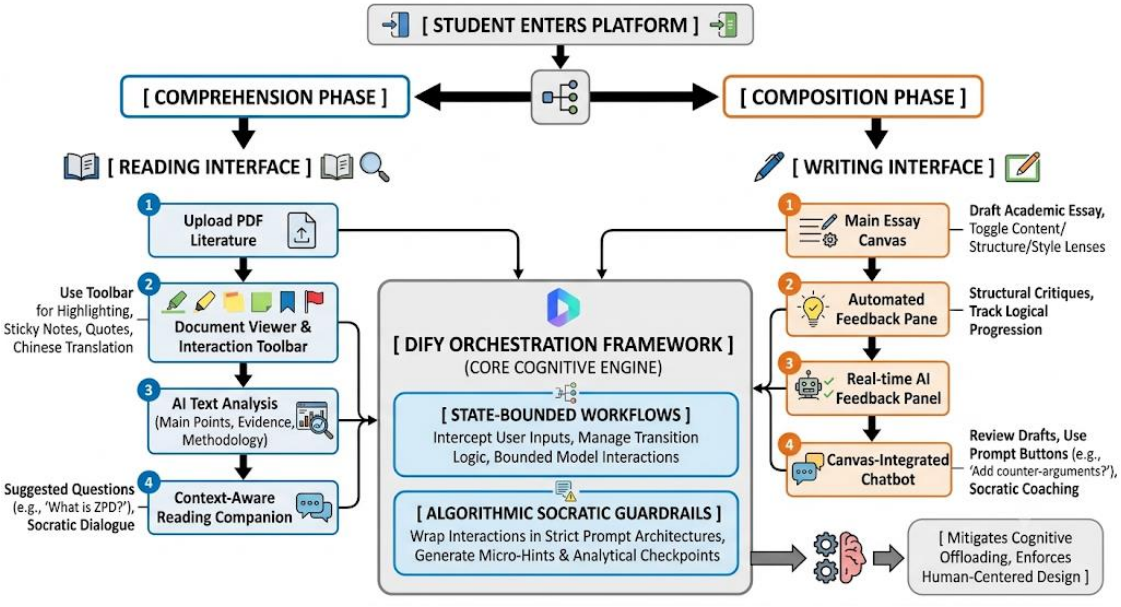


Figure 3. User workflow

The Learning Companion platform features two primary environments, the reading and writing Interfaces. It is equipped with an integrated suite of features designed for learners' needs as shown in Figure 3.

ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

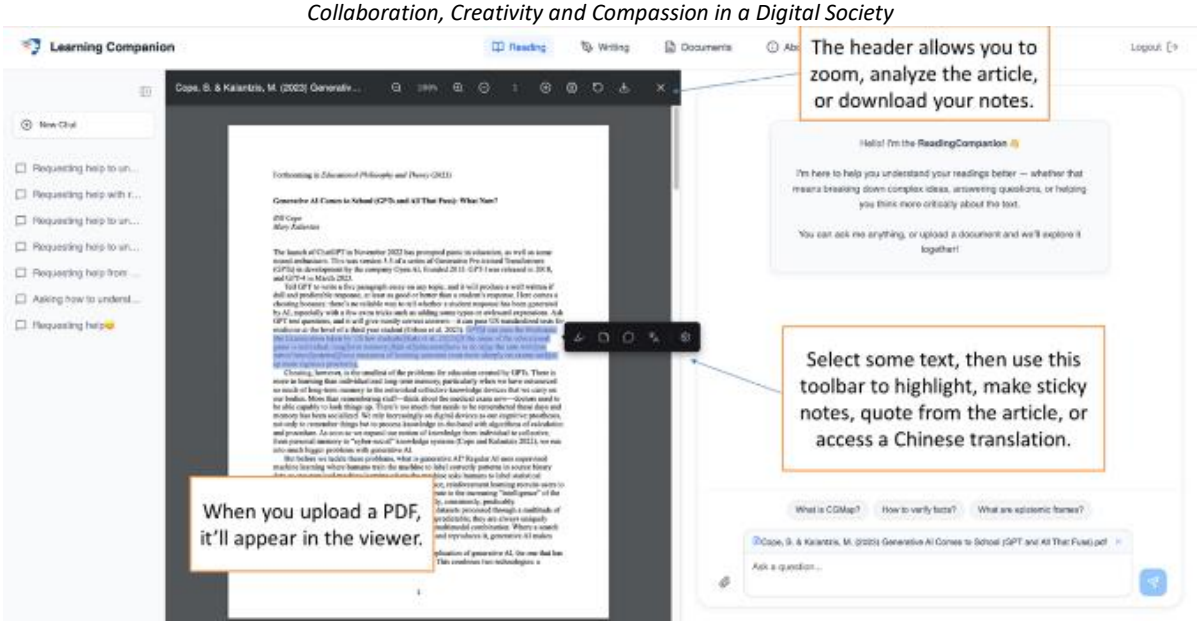


Figure 4. Reading Companion interface 1

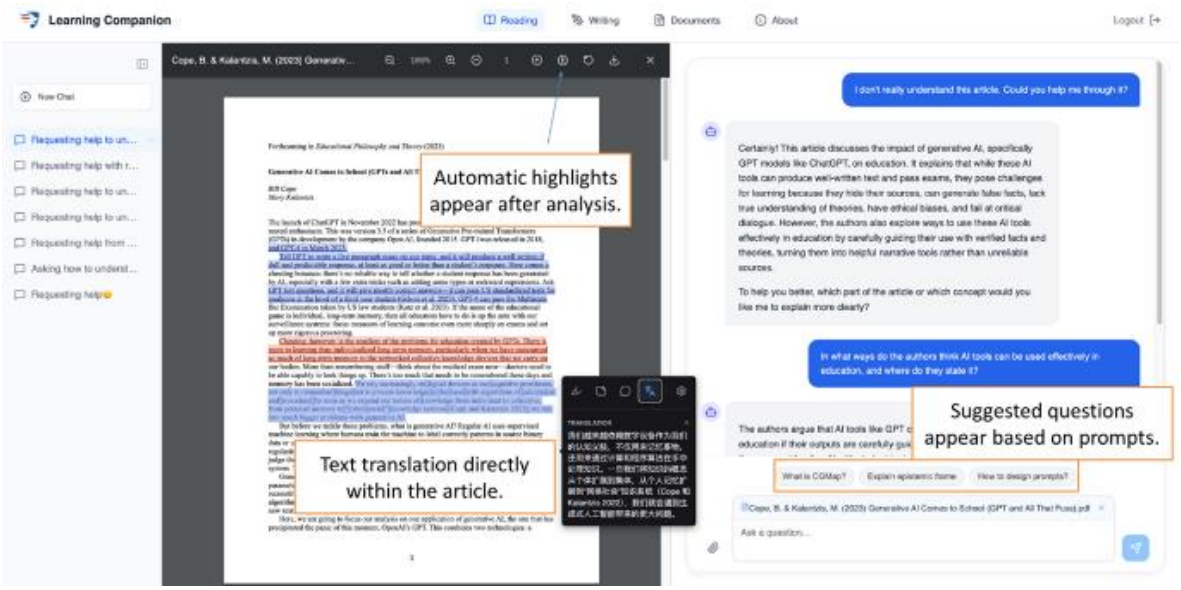


Figure 5. Reading Companion interface 2

The Reading interfaces illustrated in Figures 4 and 5 support text comprehension through three integrated components: a document viewer, an interaction toolbar, and a reading companion. Students begin by uploading the literature of their choice which is displayed as the document viewer on the left. The document viewer shows literature alongside controls for document navigation which include a component for AI text analysis. This automated analysis scans the document to identify its main points, supporting evidence, and methodology which aids reading. Students can navigate the text through the document viewer: selecting any text will trigger an interaction toolbar that contains inbuilt annotation features that include highlights, sticky notes, quotations and Chinese translation. On the right, students engage directly with the context-aware reading companion. It is designed to provide contextual guidance by asking questions or delegating small tasks which prompt students into a state of critical reasoning rather than passive consumption.

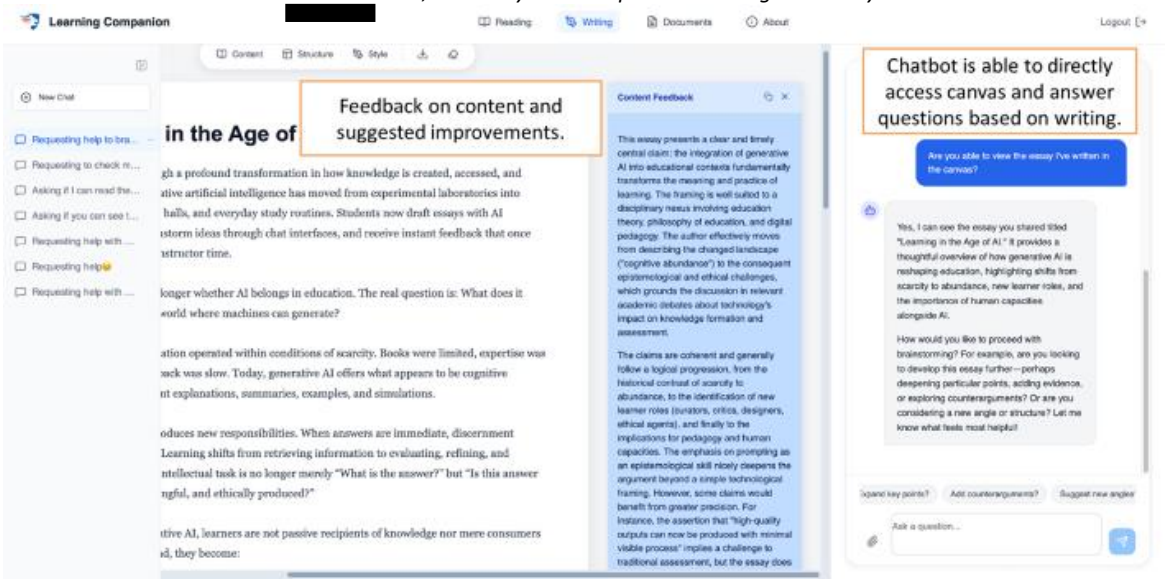


Figure 6. Writing Companion interface

The Writing interface illustrated in Figure 6 shifts the focus from comprehension to production through three core features: an essay canvas, a feedback panel, and an integrated chatbot. In the main canvas on the left, students can draft their essay while toggling between three dimensions of feedback: content, structure and style allowing students to evaluate their essays through different lenses. The feedback will be displayed within the canvas as a feedback panel. On the right, the chatbot maintains access to the canvas to evaluate drafts in real time. Rather than generating text for the student, it provides dynamic suggestion tips (e.g., “Add counter arguments?”) to push students to expand their thesis and refine their own reasoning. When the students feel unsure on how to develop their essay, they could also use the writing companion to brainstorm. Similarly, rather than generating answers, the companion keeps the cognitive effect with the student. It prompts them to clarify their thesis by, for example, identifying missing pieces within their essay.

Workflow with Dify

The system leverages Dify to separate algorithmic AI reasoning from foundational application code. Dify functions as an LLM orchestration platform: a third-party application that manages and visualises how AI models interact with prompts and inputs. This allows developers to design and manage multi-agent AI applications using graph-centric workflows and programmable nodes (LangGenius, 2024). Instead of routing user messages through a single linear text generation loop, incoming requests are passed to an intentional state machine designed around specialised task agent nodes, structuring the workflow into distinct, sequential execution stages, as illustrated in Figure 7.

The system utilizes recognized prompt engineering techniques, specifically in-context learning, role decomposition and rule-based constraints as categories in the Prompt Report systematic taxonomy (Schulhoff et al., 2024). By mapping this established prompting methodology directly into the workflow, the platform can enforce targeted instructional logic at each phase of the workflow.

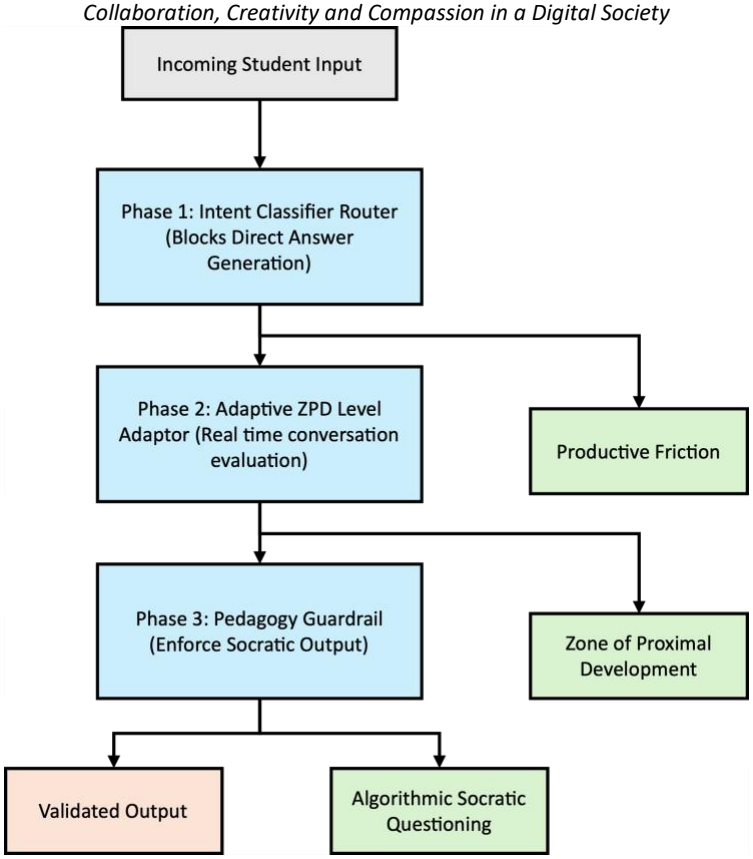


Figure 7. Illustration of Learning Companion's algorithmic workflow

Phase 1: Intent classifier

Prompt engineering method: Few-shot prompting and prompt-based rules.

Pedagogical goal (productive friction): By figuring out exactly what academic task the student is trying to tackle, the system can give highly targeted guidance instead of just generic advice. For instance, it recognizes whether a student is struggling to picture the outline or trying to address counter arguments. Regardless of which path the student takes, the underlying instructions forbid the AI from writing the actual text or giving away the answers. Instead, it forces the system to respond with conceptual frameworks or guiding hints, keeping the learning active and ensuring the student still does the heavy lifting right from the start.

Phase 2: Adaptive ZPD adaptor

Prompt engineering method: Role decomposition

Pedagogical goal (ZPD): This architecture algorithmically simulates Vygotsky's ZPD (1978). By modulating the prompt complexity and instructions dynamically session-by-session, the system ensures the assistance remains within the student's optimal cognitive threshold, preventing frustration in lower tiers and cognitive boredom or complacency in higher tiers. The system operates under the structural assumption that a student's cognitive baseline is dynamically determined in real-time based entirely on their interaction patterns within the current active conversation history. The agent runs an evaluation sweep across the current session's messaging exchange to gauge user proficiency, semantic depth, and response autonomy, subsequently tagging them into an immediate tier: Beginner, Intermediate, or Advanced, as shown in Table 1. The orchestrator then performs a runtime context injection, swapping out the active LLM instructions based on the freshly assigned conversational tier.

Table 1
Agent behaviour in relation to assigned conversational tier

Proficiency Tier	Agent Behaviour
Beginner	Restricts instructions to foundational concepts, breaking down complex terminology and defining basic structures.
Intermediate	Directs the agent to identify and challenge logical jumps in the student's text, utilizing targeted, analytical clarification questions to guide them toward tighter reasoning.
Advanced	Commands the model to assume an adversarial posture, introducing complex counter arguments and challenging core assumptions to test argumentative robustness

Phase 3: Pedagogy Guardrail

Prompt engineering method: Rule based prompts

Pedagogical goal (algorithmic Socratic questioning): By structurally restricting the output to non-directive prompts, this guardrail forces the student to remain active, preventing cognitive offloading and the illusion of competence typical of unconstrained AI interactions. It also ensures that replies are grounded by the information available

Evaluation and discussion

A 60-minute experimental study was conducted with 15 postgraduate master's students to measure system accessibility, usefulness, and meta-cognitive impact. Evaluation data was collected via a 55-question Likert-scale instrument across seven core dimensions as well as five optional open-ended questions.

Survey result

The pilot study yielded a strong validation of the system's overarching goals, as shown in Table 2. The Overall Feedback construct achieved the highest evaluation metric ($\mu = 4.16$, 73% agreement), indicating a high level of student satisfaction with the platform during the learning process. Crucially, the Learning Companion and Other GenAI construct ($\mu = 4.04$, 53% agreement) demonstrated that participants explicitly perceived this Socratic, scaffolded approach as more beneficial for educational development than consumer-grade GenAI tools focused on direct answer generation. This directly validates the core design choice to programmatically restrict cognitive offloading. The percentage agreement appears to be low throughout the construct. This is due to the aggregation within constructs for each participant as each participant may experience different pain points. On the item level, only two of the 55 questions achieve a percentage below 50%.

Table 2
Quantitative breakdown of the pilot survey results sorted by construct means and student agreement rates.

Construct	Means	Standard Deviation	Minimum	Maximum	%Agree
Learning Companion Accessibility and User Experience	3.51	0.71	1.57	4.57	20%
Overall Usefulness for learning	3.91	0.67	2.56	5.00	47%
Reading Companion Features	3.90	0.60	2.69	5.00	33%
Writing Companion Features	4.04	0.67	2.73	5.00	67%

Higher-order Learning & Metacognition	3.83	0.74	2.14	5.00	53%
Learning Companion & other GenAI	4.03	0.70	2.80	5.00	53%
Overall Feedback	4.16	0.92	2.00	5.00	73%

Pedagogy Mapping

The evaluation confirmed the system's capacity to drive higher-order cognition across its specific feature modules:

- **Writing Companion success:** The Writing Companion Feature construct performed remarkably well (μ = 4.04, 67% agreement). Notably, the item assessing whether the system "encouraged me to rethink parts of my argument" secured a mean of 4.13 (87% agreement), directly confirming that the agent's state-machine prompt logic successfully triggers critical reasoning and iterative revision.

- **Reading Companion utility:** The Reading Companion Feature achieved moderate overall utility (μ = 3.90, 33% agreement), yet its PDF-integrated document analysis feature stood out prominently, with students noting that direct canvas highlights altered how they approached reading text (μ = 4.20, 87% agreement).

- **Higher-order learning and metacognition**: Reflecting its primary educational purpose, the system achieved balanced, moderate effectiveness in promoting critical reasoning across both modules (μ = 3.83, 53% agreement), signalling a uniform scaffolding baseline but highlighting room for future framework optimisation.

Qualitative Survey

Analysis of the open-ended text responses strongly reinforces these pedagogical metrics by confirming that the system successfully enforces cognitive ownership over text production. When evaluated against standard consumer-grade AI platforms, students explicitly noted that the chatbot "doesn't give me the answers directly" and "doesn't write something for me directly," validating that the Socratic guardrails successfully disrupt passive text generation habits. Instead, the agent's tendency to pose contextual questions successfully "lead[s] me to have my own thoughts" and pushes writers to figure out "what kind of details I can add to improve this essay." The system's purposeful Socratic constraints force students into a productive struggle that handles individual baseline learning differences unevenly, while early technical friction—such as narrow chat frames and occasional application timeouts—highlights clear structural targets for future framework optimisation.

Conclusion and future work

The prototype demonstrates that when generative model interactions are bounded by multi-agent state workflows, technology can serve as a scaffold for active learning rather than an incentive for cognitive offloading. By replacing direct answer outputs with automated Socratic dialogue, the platform successfully balances institutional time limitations with human-centred pedagogy. This human-centred pedagogy is critical in the education sector as the rapid integration of GenAI risks shifting a student's role from active, critical thinker into a passive consumer of GenAI's output. By applying Socratic dialogue, the system scaffolds and mimics the productive struggle students encounter. This encourages learners to refine their own reasoning preserving human agency within learning in the new age of AI. More broadly, the findings demonstrate that AI can strengthen, rather than replace, meaningful learning relationships by supporting productive dialogue between students, educators, and intelligent learning technologies. This reinforces the importance of designing AI systems that act as collaborative learning partners, promoting critical reasoning, learner autonomy, and sustained engagement in increasingly digital educational environments.

Future work includes two technical directions. Firstly, implementing real-time adaptivity where prompt architectures, micro-hints, and question complexity automatically recalibrate based on continuous tracking of the user's reasoning proficiency. Secondly, the research can move beyond short-term user perception metrics

to establish longitudinal evaluation frameworks. This will track long-term system usage against measurable milestones to quantify the application's direct impact on sustained human cognitive improvement and academic writing development. Future research will also investigate how human-centred AI can further strengthen collaboration between students and educators while supporting reflective learning and meaningful educational interactions over time.

Acknowledgement

This study was funded by the National Institute of Education (NIE), Nanyang Technological University (NTU), Singapore, through the NIE Incentivising ICT-Use Innovations Grant (I3G; Grant No. I3G 25/25/VLF) and administered by NIE. Any opinions, conclusions, or implications in this material are those of the authors and do not necessarily reflect the views of NIE. The study was approved by the NTU Institutional Review Board (IRB) (Approval No. IRB-2025-752).

References

Abtahi, Y., Graven, M., & Lerman, S. (2017). Conceptualising the more knowledgeable other within a multi-directional ZPD. *Educational Studies in Mathematics*, 96(3), 275–287. https://doi.org/10.1007/s10649-017-9768-1

Abubakar, S., Jeilani, A., & Yusuf, M. (2025). The role of over-reliance on AI in the negative consequences of student learning: The moderating effects of ethical concerns and institutional policies. *Cogent Education, 12*(1), Article 2591503. https://doi.org/10.1080/2331186X.2025.2591503

Chodorow, K. (2013). *MongoDB: The definitive guide* (2nd ed.). O'Reilly Media.

Chukwuere, J. E. (2024). Critical research thinking: A recipe for academic writing success and publications. *Qeios*. https://doi.org/10.32388/KNWPUJ

Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education, 20*, 22. https://doi.org/10.1186/s41239-023-00392-8

Debets, T., Banihashem, S. K., Joosten-Ten Brinke, D., Vos, T. E. J., Maillette de Buy Wenniger, G., & Camp, G. (2025). Chatbots in education: A systematic review of objectives, underlying technology and theory, evaluation criteria, and impacts. *Computers & Education, 234*, 105323. https://doi.org/10.1016/j.compedu.2025.105323

Dhameria, V., Muazeib, A. I. M., Blhaj, K. M. S., Sugiyarsih, S., & Rosadah, R. A. (2025). The impact of digital transformation in higher education management: Integrating online learning and educational applications for efficiency and accessibility. *International Journal of Educational Qualitative Quantitative Research, 4*(1), 15–24. https://doi.org/10.58418/ijeqqr.v4i1.135

Favero, L., Pérez-Ortiz, Kaser, T. A., & Oliver, N. (2024). Enhancing critical thinking in education by means of a Socratic chatbot. *arXiv preprint arXiv*:2409.05511. https://doi.org/10.48550/arXiv.2409.05511

Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies, 15*(1), 6. https://doi.org/10.3390/soc15010006

Jose, B., Cherian, J., Verghis, A. M., Varghise, S. M., S, M., & Joseph, S. (2025). The cognitive paradox of AI in education: between enhancement and erosion. *Frontiers in Psychology*, 16, Article 1550621. https://doi.org/10.3389/fpsyg.2025.1550621

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., . . . Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences, 103*, 102274. https://doi.org/10.1016/j.lindif.2023.102274

LangGenius, Inc. (2024). *Dify: An open-source LLM app development platform* [Computer software]. https://github.com/langgenius/dify

Peláez, C. A., Solano, A., Núñez V, J. M., Castro, D., Cardona, J. J., Duque, J. S., Espinosa, J. C., Montaño, A. S., & Prieta, F. D. L. (2025). Designing User Experience in the Context of Human-Centered AI and Generative Artificial Intelligence: A Systematic Review. In R. Chinthaginjala, P. Sitek, N. Min-Allah, K. Matsui, S. Ossowski, S. Rodríguez (Eds.), *Distributed computing and artificial intelligence, 21st International Conference* (pp. 201–209). Springer. https://doi.org/10.1007/978-3-031-82073-1_20

Schulhoff, S., Ilie, D., Balepur, N., Kahadze, K., Liu, A., Zhou, C., Geiping, J., Carpuat, M., Goldstein, T., Shmueli, A., Small, F., Arora, S., Axelsen, A., Alon, N., Anugrah, R. S., Agnihotri, S., Alawneh, A., Atashgahi, H., Al-Khamissi, H., ... Regardie, M. (2024). The prompt report: A systematic survey of prompt engineering techniques. *arXiv preprint arXiv*:2406.06608. https://doi.org/10.48550/arXiv.2406.06608

Sharples, M. (2023). Towards social generative AI for education: Theory, practices and ethics. *Learning: Research & Practice*, *9*(2), 159–167. (173226517). https://doi.org/10.1080/23735082.2023.2261131

Tang, K.-S., & Putra, G. B. S. (2025). Generative AI as a dialogic partner: Enhancing multiple perspectives, reasoning, and argumentation in science education with customized chatbots. *Journal of Science Education and Technology*. https://doi.org/10.1007/s10956-025-10240-1

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.

Yang, Q., Yunus, M. M., & Rafiq, K. R. M. (2024). Digital innovations in higher education: A systematic review. *International Journal of Academic Research in Business and Social Sciences*, 14(2), 1330 –1344. https://doi.org/10.6007/IJARBSS/v14-i2/20630

Tan, C. Y., Fernando, O. N. N., Tan, J. H. M., Ong, C. A., Merritt, T. R., & Lim, V. F (2026). Learning companion for students' critical reasoning academic writing. In Cochrane, T., Narayan, V., Poole, T., Lai, P., Lake, P., Moore, O., & Li, L. (Eds.), *The Future is Connection: Collaboration, Creativity and Compassion in a Digital Society*. Proceedings ASCILITE 2026. Brisbane (pp. 376–386). https://doi.org/10.65106/apubs.2026.3364

Note: All published papers are refereed, having undergone a double-blind peer-review process.
The author(s) assign a Creative Commons by attribution license enabling others to distribute, remix, tweak, and build upon their work, even commercially, as long as credit is given to the author(s) for the original creation.

© Tan, C. Y., Fernando, O. N. N.., Tan, J. H. M., Ong, C. A., Merritt, T. R., & Lim, V. F. 2026