Fostering Critical AI Literacy in UK–China Transnational Computing Education: An Implementation of the SAGE Framework

Maged Fakirah
Oxford Brookes College, Chengdu University of Technology
Chengdu, Sichuan, China
maged@zy.cdut.edu.cn

Mahmoud Elkhodr
School of Engineering and Technology, Central Queensland University
Sydney, NSW, Australia
m.elkhodr@cqu.edu.au

Sajjad Hussain*
Glasgow College, University of Electronic Science and Technology of China
Chengdu, Sichuan, China
sajjad.hussain.2@glasgow.ac.uk

Abstract

As Generative AI (GenAI) becomes increasingly embedded in higher education, supporting students to engage with it critically is no longer optional, particularly in computing education, where students frequently use AI-supported tools for coding, explanation, and technical problem solving. Building on prior empirical implementations of the Structured AI-Guided Education (SAGE) framework across Australian computing-education contexts, this study examines whether SAGE can be transferred into a UK–China transnational education (TNE) setting. The intervention was conducted in a Level 4 computing tutorial and centred on a graph-theory activity involving the Minimum Spanning Tree (MST) algorithm. Rather than asking students to passively consume AI-generated answers, the activity required them to construct their own solution, critique an AI-generated response containing non-optimal edge selections, justify their corrections, adapt their solution under an additional contextual constraint, and reflect on responsible AI use. A total of 131 undergraduate computing students participated in the intervention and completed a post-activity survey. Students reported a statistically significant increase in MST confidence after the activity, with a large effect size, alongside positive perceptions of AI critique and increased caution toward unverified AI-generated technical solutions. While limited in scope to a single institution and tutorial activity in this initial TNE deployment, the study contributes early empirical evidence on critical AI literacy in UK–China transnational computing education and extends the SAGE evidence base into a new institutional context.

CCS Concepts

• **Applied computing → Interactive learning environments**; • **Theory of computation → Design and analysis of algorithms**; • **Human-centered computing →** *Empirical studies in HCI*.

*Correspondence: sajjad.hussain.2@glasgow.ac.uk

ICAIES 2026, Chengdu, China

ACM ISBN 979-8-4007-2526-5
https://doi.org/XXXXXXX.XXXXXXX

Keywords

Generative AI in Education, AI-Assisted Problem Solving, Critical AI Literacy, SAGE Framework, Computing Education, Algorithm Analysis and Design, Educational Intervention, Human-AI Interaction, Transnational Education

ACM Reference Format:
Maged Fakirah, Mahmoud Elkhodr, and Sajjad Hussain. 2026. Fostering Critical AI Literacy in UK–China Transnational Computing Education: An Implementation of the SAGE Framework. In *Proceedings of (ICAIES 2026)*. ACM, New York, NY, USA, 8 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

In recent years, Generative Artificial Intelligence (GenAI) has become increasingly embedded in higher education. Existing research describes large language models (LLMs) as offering opportunities for personalised feedback, explanation, and learning support while also raising concerns about accuracy, ethics, and over-reliance [6, 10]. In computing education in particular, students often view GenAI positively and recognise its usefulness, with prior ICT and data-oriented higher education studies reporting ChatGPT as a useful learning resource and a source of responsive learning support [5, 9].

As a result, GenAI tools have become part of students' everyday learning practices, including coding assistance, concept explanation, summarisation, brainstorming, and technical problem solving. However, although these tools offer substantial opportunities for rapid support, they also create important pedagogical concerns regarding automation bias, shallow understanding, and declining verification behaviours [2, 11]. This issue aligns with broader work on AI literacy, which is defined as the capacity to understand, use, communicate with, and critically evaluate AI technologies [7]. In other words, learners need to interrogate AI's role in knowledge production rather than merely learn how to operate such tools efficiently [8].

To address this challenge, the Structured AI-Guided Education (SAGE) framework introduced a pedagogical model centred on staged AI engagement, critical evaluation, contextual reasoning, and reflective synthesis. The framework was originally developed in cybersecurity education and was later extended to systems analysis and design contexts [3, 4]. Prior SAGE implementations demonstrated that structured orchestration activities can encourage students to move beyond passive AI consumption toward critique, contextual refinement, and analytical reasoning. In particular, the framework requires students to defend their decisions, identify AI limitations, and justify modifications rather than submit generated answers alone.

However, prior SAGE implementations were conducted primarily in Australian higher-education contexts. The transferability of the framework to transnational education (TNE) settings remains comparatively underexplored. This gap is important because TNE environments introduce distinct educational dynamics. Recent work on GenAI policy in transnational and English-medium higher education argues that institutional responses should account for contextual, cultural, linguistic, and academic-integrity differences rather than assume a single universal policy model [1]. In this context, relevant differences include Chinese–English academic language demands, diverse educational expectations, varying levels of prior exposure to AI tools, compressed tutorial structures common in TNE programmes, and different attitudes toward authority, technology, and assessment.

Accordingly, this study presents an initial implementation of SAGE within a UK–China TNE computing programme. Rather than reproducing a large multi-week deployment, the study examines whether a single-session, tutorial-scale SAGE intervention can operate within the time constraints of a normal computing tutorial. The activity centred on a graph-theory problem involving the Minimum Spanning Tree (MST) algorithm. The aim was not simply to improve algorithmic confidence, but to foster critical AI literacy through structured human–AI comparison, justification, and adaptation. This study therefore examines whether SAGE, previously established across Australian computing-education contexts, can be transferred into the distinct pedagogical, linguistic, and institutional conditions of UK–China transnational education.

The study addresses the following research questions:

- **RQ1:** Can a SAGE-guided intervention improve students' confidence in applying MST reasoning in this TNE setting?
- **RQ2:** How do students in this TNE context perceive AI critique and justification in technical problem solving, and what patterns of critical AI engagement emerge from their open-ended responses?
- **RQ3:** Does SAGE transfer effectively from its established Australian higher-education context into a UK–China transnational computing education programme?

The remainder of this paper is organised as follows. Section 2 positions the study in relation to prior SAGE implementations, critical AI literacy, and GenAI in transnational and English-medium higher education. Section 3 describes the UK–China TNE context, SAGE-guided MST tutorial design, survey instrument, classroom observation component, and analysis methods. Section 4 presents the quantitative, open-ended, and observational findings. Section 6 interprets these findings and discusses implications for TNE computing education, teacher-led GenAI classroom innovation, institutional AI literacy strategies, algorithmic pedagogy, and the SAGE research programme, before outlining limitations and future work. Section 7 concludes the paper by summarising the contribution and pointing toward graded follow-up work that implements Defend.

2 Related Work

Three genres of literature inform the present study: prior implementations of the SAGE framework, foundational works on critical AI literacy in computing education, and emerging scholarship on GenAI use within transnational and English-medium higher education.

2.1 SAGE and Structured GenAI Pedagogy

Prior SAGE implementations have established the framework's pedagogical value within Australian higher-education contexts. Initial work in cybersecurity education demonstrated that staged AI engagement, critical evaluation, and reflective synthesis can support both conceptual understanding and responsible AI use [3]. Subsequent extension into systems analysis and design provided further empirical evidence of the framework's applicability across computing sub-disciplines and supported a roadmap for embedding GenAI within technical curriculum [4]. Complementary work on ICT students' perceptions of ChatGPT and related GenAI tools informed the framework's design assumptions by identifying the opportunities and verification gaps associated with AI-supported learning [5, 9]. Across the broader SAGE research programme, the evidence base has grown beyond a single local implementation, spanning multiple empirical studies, disciplinary settings, and student cohorts. The present study extends that programme by examining whether SAGE can operate in a UK–China TNE context.

2.2 Critical AI Literacy in Computing Education

Within the broader literature on critical AI literacy, the importance of moving learners beyond operational fluency toward critical evaluation has been increasingly emphasised. Long and Magerko define AI literacy as a set of competencies that enable individuals to critically evaluate, communicate with, and use AI technologies effectively [7]. More recent scholarship on critical GenAI literacy argues that such literacy cannot be reduced to technical skill acquisition, because learners must also examine the epistemic, ethical, and relational conditions under which AI-generated outputs are produced and used [8]. These arguments are especially relevant in computing education, where students may be technically confident yet still vulnerable to over-reliance on generated explanations, code fragments, or algorithmic solutions. The SAGE framework operationalises critical AI literacy through a teachable cycle of engagement, evaluation, refinement, audit, and reflection.

2.3 GenAI in Transnational and English-Medium Higher Education

Scholarship on GenAI in TNE and English-medium higher education contexts remains comparatively thin. Existing GenAI literature has documented broad opportunities and risks in higher education [2, 6, 10], while over-reliance research has highlighted the risk that students may accept AI-generated recommendations without sufficient critical scrutiny [11]. However, these concerns may be intensified or reshaped in TNE settings, where institutional policy, curriculum expectations, English-medium instruction, and local academic cultures intersect. Bannister et al. argue that GenAI policy development in English-medium transnational higher education should account for contextual and cultural differences rather than impose universal assumptions [1]. To date, empirical investigation of structured AI-literacy pedagogies within UK–China TNE computing programmes specifically has been limited. The present study addresses this gap by testing whether a framework established in Australian computing education can be transferred into a UK–China TNE tutorial environment.

3 Methodology

This section describes the study context, intervention design, survey instrument, classroom observation component, data collection procedure, and analysis methods used to evaluate the SAGE-guided MST activity.

3.1 Educational Context

The intervention was conducted within a UK–China transnational education computing programme during a Mathematics for Computing tutorial session. It was implemented with a Level 4 cohort and focused on the MST algorithm. The tutorial was limited to 50 minutes to align with the standard teaching timetable. Accordingly, the intervention was designed as a single-session, tutorial-scale implementation rather than an extended multi-week orchestration.

3.2 SAGE-Guided Activity Design

The activity translated core SAGE principles into a five-stage tutorial workflow: (1) Human Baseline, (2) AI Solution Review, (3) AI Solution Evaluation, (4) Synthesis Under Constraint, and (5) Reflection. These stages operationalise the SAGE student cycle of Generate, Evaluate, Refine, Audit, and Reflect within a graph-theory tutorial setting. The sixth step of the educator-facing SAGE framework, Defend, is a supervised assurance step applied within graded assessment contexts. The present intervention was a formative, non-graded tutorial activity, and Defend was therefore not implemented in this study. This boundary is indicated in Fig. 1. The design positioned AI output as an object for critique rather than as a source of final answers.

The five-stage tutorial activity as shown in Figure 1 in the first five columns are mapped onto the corresponding steps of the SAGE student cycle (Generate, Evaluate, Refine, Audit, and Reflect), with the differentiated student action and role of AI indicated at each stage. The sixth step of the educator-facing SAGE framework, Defend, provides supervised assurance under graded conditions and was therefore outside the scope of the present formative, non-graded tutorial implementation; it is shown in the rightmost column for completeness. Students first constructed their own MST solution using a distance table connecting six different locations. They then examined a plausible AI-generated MST response containing multiple non-optimal edge selections. Students were subsequently asked to:

(a) identify incorrect AI decisions;
(b) justify why the selections violated MST principles;
(c) propose improved edge replacements; and
(d) explain why the corrected network satisfied MST conditions.

Finally, students modified their solution under an additional contextual constraint in which a specific location was considered highly congested and therefore should be avoided where possible. The activity was designed to move students beyond answer production. By first solving the MST independently and then critiquing a plausible but flawed AI response, students were encouraged to move from procedural application toward evaluation and justification.

Table 1: Participant profile of the UK–China TNE computing cohort

Variable	Category	n	%
Year of study	Level 4 (10 groups)	131	100
Department	Computer Science (CS)	64	48.9
	Software Engineering (SE)	67	51.1
Sex	Male	95	72.5
	Female	36	27.5

3.3 Participants

The cohort consisted of Level 4 undergraduate students from the Computer Science (CS) and Software Engineering (SE) programmes within the UK–China TNE context, with a near-balanced distribution across the two disciplines. Participation in the survey component was voluntary and anonymous. A total of 131 students participated in the activity and completed the post-activity survey. Table 1 summarises the participant profile.

3.4 Survey Instrument and Coding

The survey collected demographic characteristics, prior AI-use frequency, prior familiarity with MST, pre- and post-confidence ratings, Likert-scale reflection items, and an open-ended AI critique response. Likert-style responses were converted into numerical values for quantitative analysis, with higher values representing greater frequency, confidence, agreement, or familiarity. Open-ended responses were coded dichotomously: 0 = incorrect, vague, or unsupported reasoning; 1 = correct identification of an AI error with acceptable MST-based justification. This binary coding was selected because many student responses were brief and because the primary concern was whether students could identify at least one plausible AI error and support that identification using MST reasoning.

3.5 Instructor Observations

In addition to the post-activity survey, the lead instructor documented classroom observations during and after the tutorial. These observations were not treated as a formal qualitative dataset. Rather, they were used as a contextual source to triangulate the survey findings and to assess the feasibility of implementing SAGE within a standard 50-minute TNE tutorial structure. The observations focused on student engagement across activity stages, visible points of uncertainty, peer discussion during AI critique, and the extent to which English-language clarification was required for technical MST terminology.

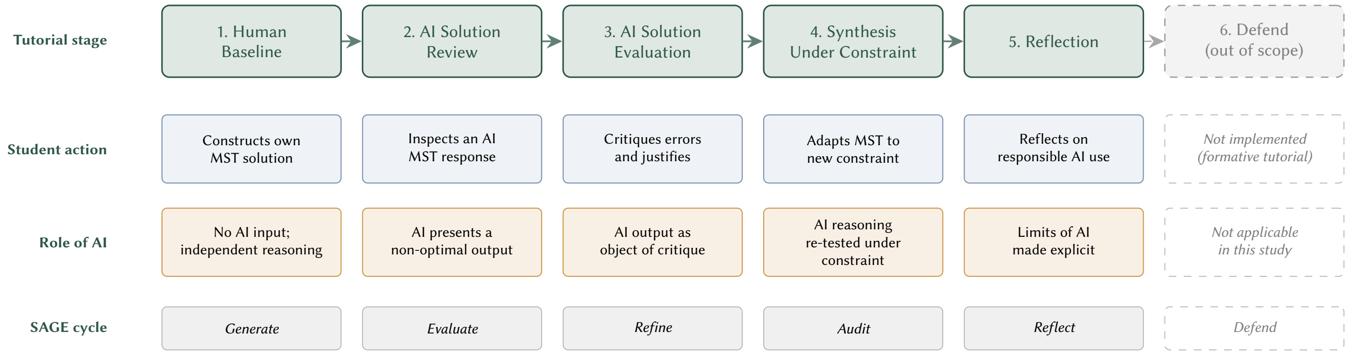

Figure 1: SAGE-guided MST tutorial workflow.

3.6 Data Analysis

Descriptive statistics were used to summarise prior AI use, MST familiarity, confidence, and post-activity perceptions. A paired-samples t-test compared students' MST confidence before and after the SAGE-guided activity. Cohen's d_z was used as the paired-samples effect-size estimate. Department-level descriptive statistics were used to examine whether the response pattern was broadly consistent across CS and SE students, but the study did not treat department as a primary inferential comparison. Open-ended responses were summarised using the binary coding scheme described above, with representative response patterns reported to illustrate the distinction between MST-based critique and unsupported or incorrect responses.

4 Results

This section presents the quantitative results, open-ended responses, and instructor observations from the SAGE-guided MST intervention.

4.1 Descriptive Profile of AI Use, Confidence, and Reflection Items

Students reported moderate AI use for general study purposes across the full cohort (M = 3.27, SD = 0.82), and AI use for technical or algorithmic problem solving was similarly common (M = 3.24, SD = 0.86). Department-level patterns are shown in Fig. 2. Overall, CS and SE students showed broadly similar patterns, although CS students reported slightly higher use on both items. These findings are consistent with the growing normalisation of GenAI tools within computing education. Informal classroom observations also suggested that many students had already used AI systems for coding support and concept explanation before the intervention.

Table 2 provides a department-level and full-cohort summary of the main survey variables. The table is descriptive and is intended to contextualise the main pre–post confidence finding and perception results rather than establish statistically significant differences between departments.

4.2 Confidence Before and After the Activity

Students showed a statistically significant increase in MST confidence following the SAGE-guided intervention. Mean confidence increased from 2.84 before the activity to 3.63 after the activity, with a large effect size. The paired-samples comparison is summarised

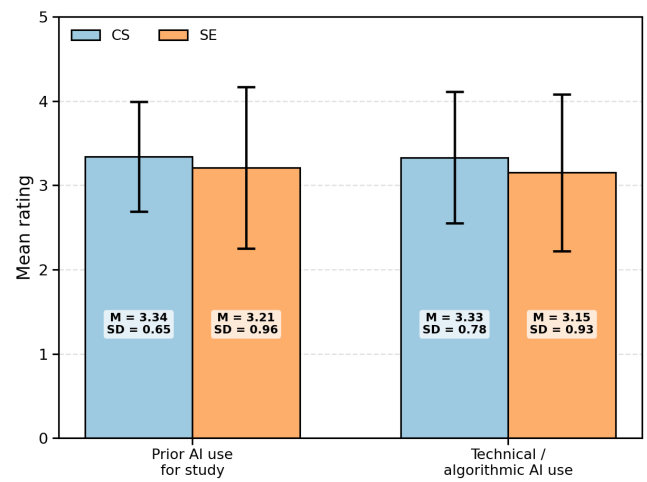

Figure 2: Prior AI use for study and technical or algorithmic problem solving by department. Error bars indicate SDs.

Table 2: Descriptive statistics for main survey variables by department and full cohort. Values are reported as $M(SD)$. Items were coded on a 1–4 scale, with higher values indicating greater frequency, confidence, familiarity, or agreement. Department-level values are descriptive only.

Variable	Total	CS	SE
Prior AI engagement			
Prior AI use for study	3.27 (0.82)	3.34 (0.65)	3.21 (0.96)
Technical/algorithmic AI use	3.24 (0.86)	3.33 (0.78)	3.15 (0.93)
Confidence and familiarity			
Algorithmic confidence before activity	3.13 (0.69)	3.14 (0.53)	3.13 (0.81)
Prior MST familiarity	1.94 (0.58)	1.88 (0.52)	2.00 (0.63)
MST confidence before activity	2.84 (0.84)	2.75 (0.78)	2.93 (0.89)
MST confidence after activity	3.63 (0.76)	3.44 (0.66)	3.81 (0.80)
Post-activity reflection items			
Justification deepened engagement	3.52 (0.64)	3.45 (0.56)	3.58 (0.70)
AI comparison helped identify errors	3.54 (0.67)	3.47 (0.62)	3.61 (0.72)
Less likely to accept AI outputs without verification	3.57 (0.77)	3.41 (0.58)	3.72 (0.90)

Table 3: Paired-samples comparison of MST confidence before and after the activity. Values are reported as $M(SD)$.

Measure	Before	After	Diff.	t(df)	p	d_z	95% CI (mean diff.)
MST confidence	2.84 (0.84)	3.63 (0.76)	0.79	10.36 (130)	$< .001$	0.91	[0.64, 0.94]

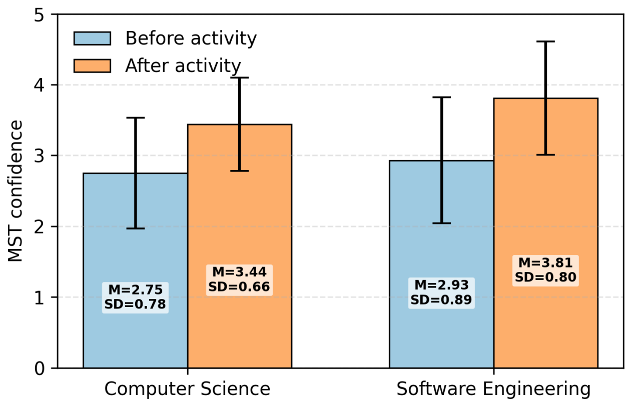

Figure 3: MST confidence before and after the SAGE-guided activity by department. Error bars indicate SDs.

in Table 3. This improvement was observed within a single standard tutorial session, suggesting that a tutorial-scale SAGE activity can produce measurable confidence gains when AI critique and justification are explicitly structured into the task.

Figure 3 shows the corresponding pre–post confidence pattern separately for CS and SE students. Both departments showed increased MST confidence after the activity, with SE students showing a slightly larger descriptive gain than CS students.

4.3 Perceptions of the SAGE Activity

Students generally responded positively to the reflective and evaluative components of the intervention. They reported that comparing their work with the AI-generated solution helped them identify mistakes that might otherwise have gone unnoticed. The strongest overall response concerned increased caution toward unverified AI outputs, with students indicating that after the activity they were less likely to accept AI-generated technical solutions without checking the underlying reasoning. These findings align with the broader SAGE objective of maintaining human authority over AI-assisted problem solving. Figure 4 summarises these perception-item means separately for CS and SE students. SE students showed slightly higher descriptive scores across the reflection items, particularly for reduced willingness to accept AI-generated technical solutions without verification.

4.4 Open-Ended Responses

The open-ended item asked students to identify one error in the AI-generated MST solution and explain why it violated MST rules. Using a strict binary coding approach, 75 of 131 responses (57.3%) were coded as correct. Correct responses identified a genuine AI

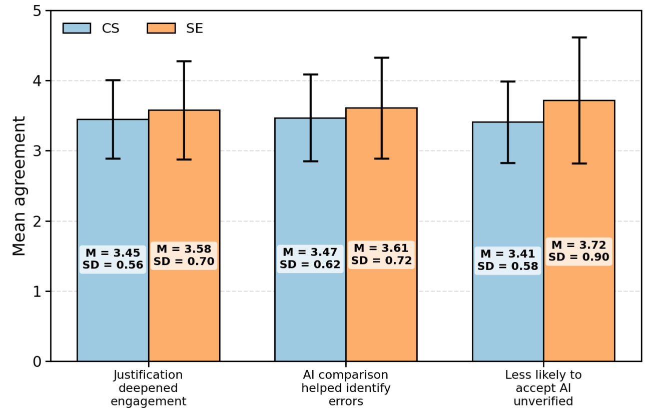

Figure 4: Students' perceptions of the SAGE-guided activity by department. Error bars indicate SDs.

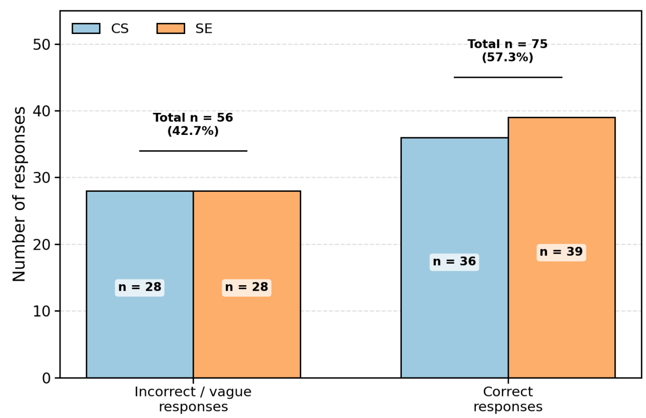

Figure 5: Binary coding outcome for open-ended MST critique responses by department.

error and linked it to valid MST reasoning, most commonly by referring to non-minimal edge selection, missed shorter valid edges, or failure to minimise the total distance. Responses that simply stated that the AI was wrong, or that incorrectly claimed that the AI solution formed a cycle, were coded as incorrect because the flawed AI solution was inefficient but did not contain a cycle. These responses suggest that the activity positioned AI output as an object of mathematical critique rather than as a source of authority. Figure 5 summarises the binary coding outcome, while Table 4 presents representative response patterns.

Descriptively, SE students produced a slightly higher proportion of correct open-ended MST critiques than CS students. As a cross-check, male and female students showed broadly similar response patterns across prior AI use, confidence change, reflection items, and open-ended MST reasoning, suggesting that the observed trends were not concentrated within a single sex group.

Table 4: Representative open-ended responses to the AI-error item

Response pattern	Example response	Coding interpretation
Specific edge-level correction	"The AI selected Galway–Sligo (85) instead of choosing the smaller valid edge Athlone–Sligo (71). Failing to choose the smallest available valid edge."	Stronger response; identifies the incorrect edge and provides a valid replacement.
Non-minimal edge selection	"Did not always select the shortest edge."	Correct MST reasoning; identifies violation of the minimum-edge principle.
Longer edge selected despite shorter alternative	"The AI selected a longer edge when a shorter one existed."	Correct MST reasoning; recognises inefficient edge selection.
General optimisation error	"It is not the minimal distance."	Partially specific but acceptable when linked to MST minimisation.
Vague or unsupported critique	"AI is wrong."	Incorrect/vague; no MST-based justification provided.
Incorrect cycle-based explanation	"The AI solution forms a cycle."	Incorrect; the flawed AI solution was inefficient but did not contain a cycle.

4.5 Classroom Observations During the Intervention

Instructor observations provided contextual evidence complementing the survey data. During the Human Baseline phase, a substantial proportion of students appeared hesitant to commit to a fully justified solution before consulting external resources, a pattern consistent with the prevalence of AI-supported study practices reported in the cohort. This hesitation diminished once the AI Solution Review phase commenced, with student engagement increasing as the AI-generated response was introduced. Across the AI Solution Evaluation phase, students engaged in peer discussion and verbal verification of edge selections, suggesting that the comparison task generated productive cognitive dissonance rather than passive consumption.

The Synthesis Under Constraint phase produced the greatest variation in observed engagement. Some groups integrated the additional contextual constraint efficiently, while others required further time to reconfigure their solutions. The Reflection phase was affected by time pressure for several groups, indicating that future deployments within 50-minute tutorial structures should allocate more time to written reflection. Across all phases, occasional English-language clarification was required for technical MST vocabulary, indicating a context-specific consideration for TNE adaptation. These observations triangulate the quantitative findings by showing that SAGE generated observable shifts in students' epistemic posture toward AI-generated content within a compact classroom activity.

5 Discussion

6 Discussion

Students' confidence in applying MST reasoning increased significantly following the single-session tutorial intervention, with a large paired-samples effect size, and the post-activity reflection items indicated positive perceptions of structured AI critique alongside increased scepticism toward unverified AI-generated technical solutions. These outcomes were observed within a UK–China TNE computing cohort operating under English-medium instruction and within a 50-minute tutorial structure. Three findings carry particular weight for AI-integrated computing pedagogy and, in turn, for the broader SAGE research programme.

First, the results indicate that critique-oriented activity supported conceptual engagement more directly than answer-focused practice alone. The open-ended responses show that many students engaged with MST reasoning principles, particularly minimum-edge selection and overall edge optimisation, rather than with the AI system itself. By positioning the AI-generated solution as an object for evaluation rather than as a source of correct answers, the activity shifted students from procedural application toward justification and verification. The instructor observations corroborate this pattern, with visible engagement increasing at the point of AI introduction and active peer verification characterising the evaluation phase. This finding is consistent with the broader argument within critical AI literacy scholarship that learners need to interrogate AI's role in knowledge production rather than only learn to operate it efficiently [8].

Second, the results indicate that measurable gains in critical AI engagement can be developed within compressed teaching structures. The intervention did not rely on prohibition or generic caution. Students were required to identify specific MST errors, to justify the violations against MST principles, and to propose corrected solutions. The department-level consistency between Computer Science and Software Engineering responses further suggests that the observed effect was not bound to a single disciplinary cohort within the programme. For TNE computing education specifically, this indicates that tutorial-level redesign offers a viable point of intervention for AI-aware pedagogy where whole-curriculum redesign is not feasible.

Third, the results indicate that structured AI-literacy pedagogy can be deployed effectively within the UK–China TNE computing setting. The statistically significant gain in MST confidence, the positive perceptions of the AI critique task, and the increased scepticism toward unverified AI outputs were each evidenced within a single 50-minute tutorial. Instructor observations further indicated that students engaged readily with the critique task, used peer discussion productively during the AI evaluation phase, and produced MST-grounded critiques in their open-ended responses. Occasional English-language clarification was required for technical MST vocabulary, which is a routine feature of English-medium computing instruction. Taken together with the quantitative results, the activity operated successfully within the standard tutorial structure of the programme without requiring substantial pedagogical redesign. This carries direct implications for cross-cultural and transnational computing programmes more broadly. The present results show that critical AI literacy pedagogy can be delivered effectively within the compressed delivery formats and shared module specifications common to TNE settings, with measurable learning indicators across confidence, perception, and critical AI engagement to support its further deployment.

6.1 Implications

The findings carry implications for transnational computing education, for teacher-led pedagogical innovation in GenAI-integrated classrooms, for institutional AI literacy strategies, for algorithmic and graph-theory pedagogy, and for the broader SAGE programme of research.

For computing education within transnational and English-medium higher-education contexts, the study indicates that critical AI literacy can be embedded within standard tutorial structures without whole-curriculum redesign. This is consequential because TNE programmes typically operate under compressed timetables and shared content frameworks that constrain the discretion available to local academics. The five-stage activity template developed within the present implementation, comprising human baseline construction, AI solution review, AI solution evaluation, synthesis under constraint, and reflection, has been shown to function within the 50-minute tutorial structure typical of UK–China TNE delivery. The cohort response, drawn from Level 4 Computer Science and Software Engineering students within a UK–China TNE programme, also contributes initial empirical evidence specific to UK–China transnational computing learners, a population for which the existing GenAI-pedagogy literature remains comparatively thin.

For teacher-led pedagogical innovation in GenAI-integrated computing classrooms, the study evidences that measurable improvements in students' critical engagement with AI can be produced by a single academic at a partner institution, working within a normal teaching allocation and without large-scale resource investment or institutional restructuring. This pathway is particularly relevant within TNE settings, where local academics frequently operate within shared module specifications and limited discretionary authority over assessment design but retain pedagogical agency at tutorial level. The implementation reported here demonstrates that tutorial-level redesign, when grounded in a structured pedagogical model, can produce statistically significant confidence gains and observable shifts in students' epistemic posture toward AI-generated content.

For institutional AI literacy strategies, the findings indicate that responsible GenAI engagement is not contingent on AI prohibition or extensive policy frameworks alone. Pedagogical design can directly shape students' epistemic posture toward AI, with structured comparison, justification, and constrained adaptation proving more educationally productive than warning-based approaches. This carries particular relevance for TNE institutions, which must reconcile differing AI-use expectations and regulatory orientations across partner jurisdictions during a period of continuing institutional uncertainty regarding GenAI in higher education.

For algorithmic and graph-theory pedagogy specifically, the study identifies a productive fit between optimisation problems and structured human–AI comparison tasks. Optimisation errors are visible, defensible, and logically explainable, which makes them well suited to critique-oriented activities. Similar problem classes, including sorting, shortest path, dynamic programming, and other classical algorithmic structures, may benefit from analogous treatments and represent a productive avenue for further tutorial-scale pedagogical research within TNE computing programmes.

For the SAGE programme of research, the study extends the framework's empirical base into a new transnational institutional context and demonstrates that the first five steps of the cycle can operate effectively within a single 50-minute tutorial without substantial adaptation, supporting a portable tutorial-scale mode alongside the more extensive course-level implementations previously reported.

6.2 Limitations and Future Work

Several limitations should be noted. The study was conducted at a single institution with a single tutorial activity, and the survey relied on self-reported confidence and perception measures rather than longitudinal performance data. Some open-ended responses were brief, indicating that gains in confidence may have exceeded gains in justificatory depth. The cohort was nationally homogeneous despite the TNE framing, and no control or comparison condition was included. The instructor observations provided useful contextual triangulation, but they were not collected through a formal ethnographic or multi-observer protocol. Because Defend was not implemented, the study cannot determine whether students could subsequently demonstrate ownership of their MST reasoning under supervised conditions; this complementary evidence is being addressed in companion SAGE work currently under development. Future work should examine whether repeated SAGE activities produce more durable AI-verification habits across different TNE contexts, institutions, and computing topics. Comparative designs may also help distinguish the effect of SAGE-guided AI critique from conventional tutorial practice.

7 Conclusion

This study presented an initial implementation of the SAGE framework within a UK–China transnational computing education environment. Using a short MST-focused tutorial intervention, students were required to construct, critique, defend, and refine AI-assisted solutions rather than passively consume generated answers. The findings revealed significant improvement in confidence, positive perceptions toward reflective critique activities, and emerging scepticism toward unverified AI-generated technical solutions. Importantly, the intervention remained operational within a normal tutorial structure, suggesting that SAGE can transfer into constrained TNE teaching environments as a practical, classroom-deployable approach to critical AI literacy.

The contribution of the study lies in providing early empirical evidence on critical AI literacy in UK–China transnational computing education, and in extending an established SAGE evidence base into a new international teaching context. Future work will include a graded follow-up study in which Defend is also implemented, completing the SAGE cycle within a single TNE deployment and allowing students' ownership of MST reasoning to be assessed under supervised conditions. A broader journal-scale investigation will involve additional TNE institutions, richer qualitative analysis, longitudinal measurements, and comparative orchestration designs. Future research should also investigate whether repeated SAGE exposure produces durable behavioural changes in students' AI verification practices beyond single-session interventions.

Acknowledgments

The authors thank the participating students for their engagement in the tutorial activity. The study was supported by the University of Glasgow.

References

[1] Peter Bannister, Elena Alcalde Peñalver, and Alexandra Santamaría Urbieta. 2024. Transnational Higher Education Cultures and Generative AI: A Nominal Group Study for Policy Development in English Medium Instruction. *Journal for Multicultural Education* 18, 1/2 (2024), 173–191. doi:10.1108/JME-10-2023-0102

[2] Cecilia Ka Yuk Chan and Wenjie Hu. 2023. Students' Voices on Generative AI: Perceptions, Benefits, and Challenges in Higher Education. *International Journal of Educational Technology in Higher Education* 20 (2023), 43. doi:10.1186/s41239-023-00411-8

[3] Mahmoud Elkhodr and Ergun Gide. 2025. The SAGE Framework for Developing Critical Thinking and Responsible Generative AI Use in Cybersecurity Education. *Discover Education* 4 (2025), 517. doi:10.1007/s44217-025-00935-3

[4] Mahmoud Elkhodr and Ergun Gide. 2026. AI Leads, Humans Lead, or Collaborate? Empirical Findings and the SAGE Roadmap for Embedding GenAI in Systems Analysis and Design Education. *STEM Education* 6, 2 (2026), 194–229. doi:10.3934/steme.2026009

[5] Mahmoud Elkhodr, Ergun Gide, Robert Wu, and Omar Darwish. 2023. ICT Students' Perceptions Towards ChatGPT: An Experimental Reflective Lab Analysis. *STEM Education* 3, 2 (2023), 70–88. doi:10.3934/steme.2023006

[6] Enkelejda Kasneci, Kathrin Sessler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günnemann, Eyke Hüllermeier, Stephan Krusche, Gitta Kutyniok, Tilman Michaeli, Claudia Nerdel, Jürgen Pfeffer, Oleksandra Poquet, Michael Sailer, Albrecht Schmidt, Tina Seidel, Matthias Stadler, Jochen Weller, Jochen Kuhn, and Gjergji Kasneci. 2023. ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education. *Learning and Individual Differences* 103 (2023), 102274. doi:10.1016/j.lindif.2023.102274

[7] Duri Long and Brian Magerko. 2020. What Is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3313831.3376727

[8] Chrysi Rapanta, Ibrar Bhatt, Aras Bozkurt, Jennifer Chubb, Christopher Erb, Ingrid Forsler, Karen Gravett, Marguerite Koole, Tim Lintner, Alex Örtegren, Tiziana Petricini, Bethan Rodgers, Jenny Webster, Xiaoxiao Xu, Inger-Marie F. Christensen, Nina Bonderup Dohn, Lene Leth W. Christensen, Sandris Zeivots, and Petar Jandrić. 2025. Critical GenAI Literacy: Postdigital Configurations. *Postdigital Science and Education* 7 (2025), 1296–1333. doi:10.1007/s42438-025-00573-w

[9] Raj Sandu, Ergun Gide, and Mahmoud Elkhodr. 2024. The Role and Impact of ChatGPT in Educational Practices: Insights from an Australian Higher Education Case Study. *Discover Education* 3 (2024), 71. doi:10.1007/s44217-024-00126-6

[10] Ahmed Tlili, Boulus Shehata, Michael Agyemang Adarkwah, Aras Bozkurt, Daniel T. Hickey, Ronghuai Huang, and Brighter Agyemang. 2023. What If the Devil Is My Guardian Angel: ChatGPT as a Case Study of Using Chatbots in Education. *Smart Learning Environments* 10, 1 (2023), 15. doi:10.1186/s40561-023-00237-x

[11] Chunpeng Zhai, Santoso Wibowo, and Lily D. Li. 2024. The Effects of Over-Reliance on AI Dialogue Systems on Students' Cognitive Abilities: A Systematic Review. *Smart Learning Environments* 11 (2024), 28. doi:10.1186/s40561-024-00316-7