Generative AI and Student Creativity: Cognitive Scaffolding Tool or Constraint on Originality? An Empirical Investigation in Vietnamese Higher Education

Nguyễn Thị Mai

Department of Educational Technology and Innovation, University of Education, Vietnam National University

Correspondence: *mai.nguyen@vnu.edu.vn*

Target Track: Emerging Technologies in Education & Social Sciences (ICCMETS 2026)

Abstract

The rapid integration of Generative Artificial Intelligence (GenAI) into higher education has sparked a profound pedagogical debate: Does generative technology expand students' creative frontiers as an interactive cognitive scaffolding tool, or does it homogenize intellectual exploration and erode originality? This study investigates the dualistic impact of large language models (LLMs) on undergraduate students' divergent thinking, conceptual originality, and domain-specific creative problem-solving within the Vietnamese higher education sector.

Employing a mixed-methods sequential explanatory design, we conducted a controlled experimental study ($N = 480$) across three major universities in Hanoi, Da Nang, and Ho Chi Minh City. Participants engaged in standardized creative concept-generation and essay-design tasks under three randomized experimental conditions: (a) Control (unaided cognitive generation), (b) Unconstrained GenAI Access (direct text generation via commercial LLMs), and (c) Structured Scaffolding (prompt-guided iterative dialogue emphasizing critique and divergence). Outputs were evaluated using the Consensual Assessment Technique (CAT) by an expert panel, complemented by Torrance Tests of Creative Thinking (TTCT) metrics (Fluency, Flexibility, Originality, and Elaboration) and computational semantic distance analysis.

Quantitative findings revealed an empirical paradox: while unconstrained GenAI access significantly increased task fluency ($F(2,477) = 48.21, p < .001, \eta^2 = .168$) and structural elaboration, it produced a statistically significant reduction in conceptual originality ($d = -0.58, p < .001$) compared to the unaided control group. Conversely, students operating under structured metacognitive scaffolding achieved superior overall creativity profiles, balancing semantic breadth with genuine novelty. Qualitative prompt-log analyses ($n = 36$) revealed widespread "early-stage design fixation" and passive cognitive offloading among unconstrained users.

This paper articulates the *Cognitive Anchor-Rebound Model* (CARM) and provides actionable instructional frameworks for educators seeking to harness AI as an intellectual collaborator rather than an algorithmic crutch.

Keywords: Generative Artificial Intelligence, Higher Education, Student Creativity, Divergent Thinking, Originality, Cognitive Scaffolding, Design Fixation, Vietnamese Higher Education.

1. Introduction

The commercialization and ubiquitous availability of Generative Artificial Intelligence (GenAI)—principally powered by Large Language Models (LLMs) such as OpenAI's GPT-series, Google's Gemini, and Anthropic's Claude—has catalyzed a paradigm shift across global higher education systems (Dwivedi et al., 2023; Kasneci et al., 2023). Within developing

knowledge economies such as Vietnam, where the Ministry of Education and Training (MOET) has aggressively prioritized tertiary digital transformation (Do et al., 2024), generative tools are increasingly woven into daily student scholarship. Undergraduates utilize LLMs for literature synthesis, software debugging, linguistic translation, and open-ended creative ideation.

However, the educational and cognitive ramifications of GenAI on higher-order thinking skills remain sharply contested. Creativity—traditionally conceptualized as the capacity to produce work that is both novel (original, unexpected) and appropriate (useful, adaptive) (Amabile, 1996; Runco & Jaeger, 2012)—occupies a central position in 21st-century educational competencies.

Proponents argue that GenAI serves as a powerful instrument of *distributed cognition* (Hutchins, 1995; Salomon, 1993), functioning as an infinitely patient brainstorm partner that lowers cognitive barriers, disrupts associative blocks, and democraticizes ideation (Baidoo-Anu & Ansah, 2023; Mollick & Mollick, 2023). Under this constructivist lens, GenAI provides adaptive cognitive scaffolding, allowing students to rapidly iterate concepts beyond their existing lexical or domain constraints.

Conversely, educational psychologists and cognitive scientists caution against premature techno-optimism. When students outsource the initial, messy phase of divergent thinking to deterministic, probabilistic algorithms, they risk succumbing to *algorithmic conformity* and *design fixation* (Jansson & Smith, 1991; Sio et al., 2015). Because LLMs generate responses derived from statistical token probabilities trained on massive historical corpora, their outputs intrinsically gravitate toward the central tendencies of conceptual spaces—yielding answers that are syntactically polished and semantically plausible, yet fundamentally unoriginal (Bender et al., 2021; Cotton et al., 2024). In academic cultures historically influenced by rote-learning paradigms, such as Vietnam's transitioning tertiary sector, excessive reliance on synthetic text engines threatens to supplant authentic critical inquiry with fluent, homogenized imitation.

Despite mounting theoretical commentary, empirical studies systematically examining the boundary conditions between AI-as-scaffold and AI-as-constraint within Southeast Asian university cohorts remain scarce. This research addresses that empirical gap by investigating the following core research questions:

1. How does the integration of Generative AI influence students' creative performance across the four standard divergent thinking dimensions: fluency, flexibility, elaboration, and originality?

2. To what extent does the structure of pedagogical intervention (unconstrained use vs. guided metacognitive scaffolding) mitigate algorithmic design fixation?

3. What cognitive patterns and prompt engineering behaviors characterize students who successfully maintain intellectual originality while leveraging generative systems?

2. Theoretical Framework & Literature Review

2.1 The Dual-Factor Model of Creativity

Creativity assessment in educational research relies predominantly on the dual-factor definition advanced by Runco and Jaeger (2012) and Amabile (1996): an artifact or idea must simultaneously satisfy criteria of *originality* (statistical infrequency, uniqueness, or novelty) and *appropriateness* (task relevance, utility, or contextual validity). In cognitive psychometrics,

divergent thinking serves as the leading indicator of creative potential, operationalized via the Torrance Tests of Creative Thinking (TTCT; Torrance, 1966):

- **Fluency:** The total quantity of interpretable, task-relevant ideas generated.
- **Flexibility:** The variety of distinct conceptual categories or conceptual shifts exhibited.
- **Originality:** The statistical rarity or non-obviousness of generated ideas relative to the peer cohort.
- **Elaboration:** The depth, detail, and developmental richness applied to an idea.

2.2 Cognitive Scaffolding, Cognitive Offloading, and Distributed Cognition

Vygotsky's (1978) *Zone of Proximal Development* (ZPD) posits that learners achieve higher developmental milestones when supported by structured environmental scaffolding. In technological environments, cognitive tools can function as cognitive partners (Salomon et al., 1991). When students engage with LLMs through iterative prompting, the system acts as an externalized semantic memory network, reducing baseline intrinsic cognitive load (Sweller, 1988) and freeing working memory capacity for higher-order synthesis and metacognitive evaluation (Mollick & Mollick, 2023).

Nonetheless, distributed cognition can degenerate into premature *cognitive offloading* (Risko & Gilbert, 2016). When the cognitive effort required to synthesize divergent concepts is completely bypassed, the neural and associative mechanisms responsible for deep structural insight remain unactivated. The student transitions from a reflexive thinker to a passive curator of pre-digested computational outputs.

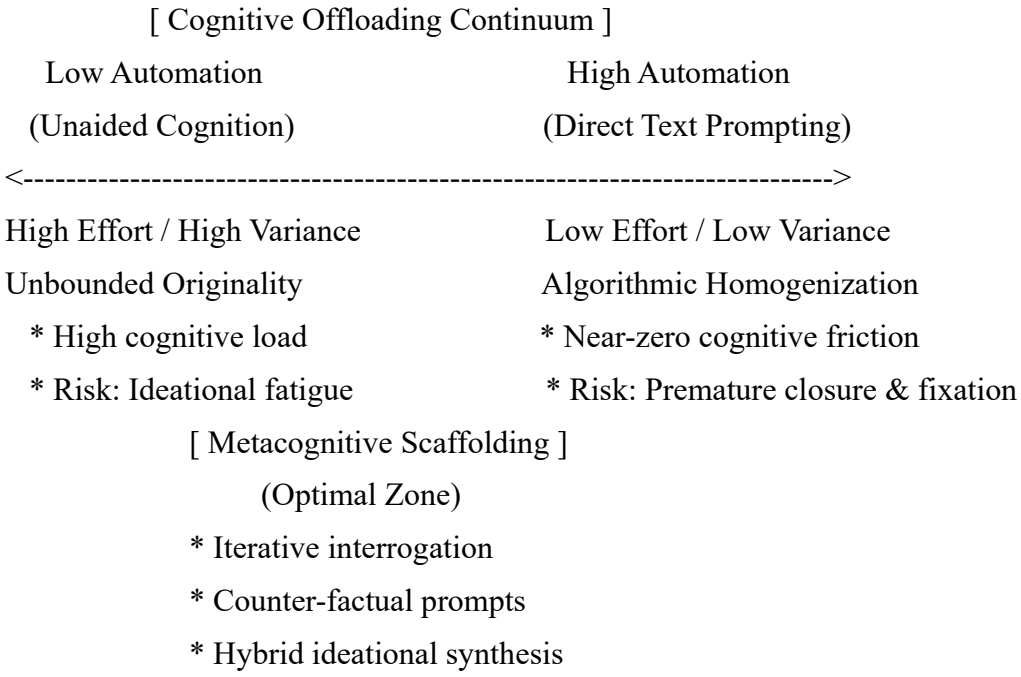

2.3 The Mechanics of Algorithmic Fixation

Design fixation occurs when an agent's creative trajectory is unconsciously circumscribed by initial exposure to salient examples or pre-existing templates (Jansson & Smith, 1991). LLMs, operating via auto-regressive next-token prediction, prioritize high-probability sequences across multidimensional vector spaces.

Consequently, AI-generated solutions reflect the statistical median of human culture. When a student reviews an initial LLM output, that synthesized text serves as a powerful cognitive anchor (Tversky & Kahneman, 1974). Divergent associative jumps become constrained within the semantic neighborhood defined by the AI's initial response.

2.4 The Higher Education Context of Vietnam

Higher education institutions in Vietnam are currently undergoing comprehensive structural reforms under Resolution No. 29-NQ/TW, shifting pedagogical emphasis from traditional teacher-centric knowledge transmission toward learner-centered constructivism and innovation (Tran et al., 2020). However, large class sizes, resource constraints, and lingering exam-oriented educational practices continue to challenge pedagogical agility.

The arrival of free and low-cost GenAI tools provides immediate utility to Vietnamese university students seeking to overcome academic writing hurdles and English-language barriers. Concurrently, Vietnamese academic leadership faces an urgent dilemma: how to prevent students from adopting GenAI as a replacement for independent conceptual exploration while equipping them with modern technical literacy.

3. Research Methodology

3.1 Experimental Design

This study utilized a mixed-methods sequential explanatory design ($QUAN \rightarrow qual$). The quantitative phase deployed a randomized between-subjects experimental design with three conditions:

1. **Condition 1: Unaided Control ($n = 160$):** Participants completed creative problem-solving and ideation tasks relying entirely on unassisted cognitive faculties. No digital devices, search engines, or AI interfaces were permitted.

2. **Condition 2: Unconstrained GenAI Access ($n = 160$):** Participants were granted unmonitored access to an LLM interface (GPT-4o / Claude 3.5 Sonnet). They were instructed: *"You may use the AI tool in any manner you choose to complete this task."*

3. **Condition 3: Structured Scaffolding ($n = 160$):** Participants interacted with the same LLM interface but were required to execute tasks following the *Cognitive Anchor-Rebound Protocol* (CARP). This protocol mandated:

   - Phase A: Independent generation of at least three raw hypotheses/ideas prior to opening the interface.
   - Phase B: Targeted AI interrogation using critique and counter-factual prompts (e.g., *"Identify five unstated assumptions and three non-obvious failure modes of my idea"*).
   - Phase C: Synthesis and divergent expansion without copying direct synthetic text blocks.

3.2 Participants and Sampling Matrix

The sample was recruited using stratified random sampling across three major Vietnamese tertiary institutions: Vietnam National University (Hanoi), The University of Danang, and Ho Chi Minh City University of Technology (HUTECH). A total of 480 undergraduate students participated across three broad disciplinary clusters.

```
Total Participants: N = 480
├── Humanities & Social Sciences: n = 160 (33.3%)
├── Business & Management:         n = 160 (33.3%)
└── Design, Media & Technology:   n = 160 (33.3%)
Gender Distribution:
├── Female: 58.3% (n = 280)
├── Male:   39.6% (n = 190)
└── Non-binary / Undisclosed: 2.1% (n = 10)
```
Mean Age: 20.4 years (SD = 1.35)

3.3 Experimental Tasks

Participants were administered two interconnected creative tasks within a 90-minute timed session:

- **Task 1: Divergent Problem Solving (Alternative Solutions Task - AST):** Adapted from Guilford's and Torrance's frameworks. Students were tasked with proposing innovative, sustainable solutions to mitigate severe urban microplastic contamination within the Mekong Delta agricultural belt.
- **Task 2: Conceptual Project Proposal:** Students authored a 600-word conceptual campaign brief for a socio-technological enterprise addressing rural education equity in Vietnam.

3.4 Measurement Instruments and Evaluation

Performance was evaluated across quantitative divergent thinking metrics, expert subjective ratings, and computational NLP vector dispersion:

1. **Divergent Thinking Coding (TTCT Criteria):** Two blinded, trained raters scored all submissions across Fluency (count of distinct, task-appropriate concepts), Flexibility (number of distinct conceptual domains utilized), and Elaboration (rated on a 1–5 rubric evaluating structural depth). Inter-rater reliability was high (Cohen's $\kappa = .88$).

2. **Consensual Assessment Technique (CAT):** Following Amabile (1996), a panel of six independent domain experts (two university professors in environmental social science, two creative design directors, and two innovation hub leads) independently evaluated the project briefs. Experts rated *Originality*, *Appropriateness*, and *Overall Creativity* on 7-point Likert scales ($1 = \text{Extremely Low}, 7 = \text{Exceptional}$). The panel demonstrated strong internal consistency (Cronbach's $\alpha = .91$).

3. **Computational Semantic Distance (Algorithmic Novelty):** To guard against subjective human bias, submissions were embedded into a 1536-dimensional vector space using OpenAI's text-embedding-3-large. For each condition, we calculated the pairwise cosine similarity among all student submissions within the same task:

$$\text{Cosine Similarity}(u, v) = \frac{u \cdot v}{\|u\|_2 \|v\|_2}$$

A higher mean cosine similarity indicates conceptual clustering (homogenization), whereas lower cosine similarity indicates greater semantic dispersion (population-level diversity).

3.5 Qualitative Data Collection

Following experimental testing, 36 participants (12 from each condition, stratified across performance quartiles) completed 30-minute semi-structured retrospective interviews.

Additionally, full digital interaction transcripts, prompt sequences, timestamped edits, and revision trees were exported from the test environment for thematic coding.

4. Quantitative Results and Statistical Analysis

4.1 Comparative Analysis Across Experimental Conditions

A one-way multivariate analysis of variance (MANOVA) was conducted to evaluate the effect of experimental condition (Control vs. Unconstrained GenAI vs. Structured Scaffolding) on the primary dependent variables: Fluency, Flexibility, Originality (CAT), and Semantic Distance.

Table 1: Descriptive Statistics and Univariate ANOVA across Conditions (N = 480)

Variable	Control (n=160) Mean (SD)	Unconstrained (n=160) Mean (SD)	Scaffolded (n=160) Mean (SD)	F(2, 477)	p-value	η^2
Fluency (Count)	5.12 (1.42)	11.45 (2.85)	8.84 (1.95)	284.15	<.001	.543
Flexibility (Domains)	3.18 (0.91)	4.22 (1.14)	5.61 (1.08)	179.42	<.001	.429
Elaboration (1-5)	2.65 (0.78)	4.15 (0.62)	4.38 (0.55)	312.60	<.001	.567
Originality (CAT 1-7)	4.82 (1.15)	3.41 (0.94)	5.44 (0.88)	165.23	<.001	.409
Appropriateness (1-7)	4.21 (1.02)	5.28 (0.75)	5.62 (0.68)	118.70	<.001	.332
Overall Creativity	4.48 (0.98)	3.95 (0.81)	5.51 (0.74)	134.88	<.001	.361

Note: Means represent composite ratings. Degrees of freedom for all univariate tests: (2, 477).

The omnibus MANOVA revealed a statistically significant difference among conditions across the combined dependent measures:

$$\text{Pillai's Trace} = 0.941, \quad F(12,946) = 68.34, \quad p < .001, \quad \eta_p^2 = .470$$

Subsequent univariate testing (Table 1) and post-hoc pairwise Tukey HSD comparisons revealed distinct dynamics:

- **The Fluency-Originality Divergence:** Participants in the Unconstrained GenAI condition generated more than twice as many raw ideas as unaided students ($M = 11.45$ vs. $M = 5.12, p < .001$). However, their expert-rated Originality scores plummeted ($M = 3.41$ vs. $M = 4.82, p < .001, d = -1.34$). Unassisted students produced fewer concepts, but those concepts demonstrated significantly higher originality and conceptual non-obviousness.
- **The Superiority of Structured Scaffolding:** Participants in the Scaffolded condition exhibited the most balanced creative profile. They achieved significantly higher Originality ($M = 5.44$) than both the Control ($p < .001, d = 0.61$) and Unconstrained groups ($p < .001, d = 2.23$). Furthermore, their conceptual Flexibility (number of distinct domains) exceeded all other groups ($M = 5.61, p < .001$).

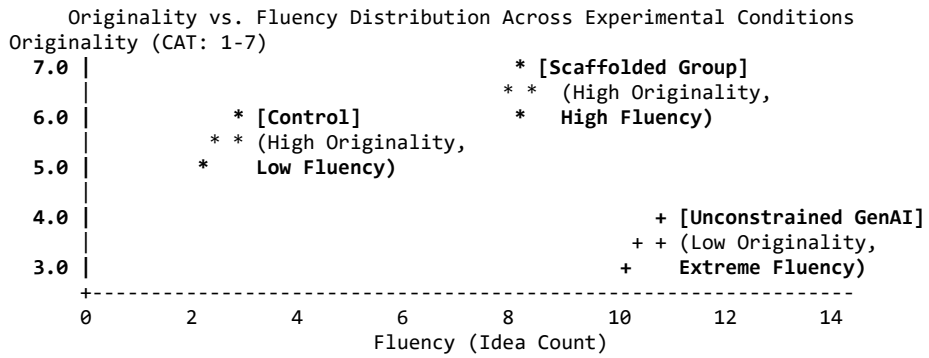

4.2 Computational Semantic Dispersion and Homogenization

Analysis of the 1536-dimensional embedding vectors provided objective mathematical confirmation of algorithmic homogenization.

Table 2: Intra-Group Cosine Similarity and Semantic Dispersion Matrix

Condition	Mean Pairwise Cosine Similarity (SD)	Semantic Dispersion Index (1 - CosSim)
Unaided Control	0.412 (0.084)	0.588
Unconstrained GenAI	0.689 (0.061)	0.311
Structured Scaffolding	0.465 (0.072)	0.535

The mean pairwise cosine similarity between student submissions in the Unconstrained GenAI condition was $\rho = 0.689$, compared to $\rho = 0.412$in the Unaided Control group ($t(318) = 34.21, p < .001, d = 3.79$).

Submissions in the Unconstrained condition were tightly clustered within a narrow semantic subspace. Unassisted students, while slower, explored geographically disparate sectors of the conceptual landscape. The Scaffolded group maintained high semantic dispersion ($\rho = 0.465$,Dispersion $= 0.535$), effectively combining the expansive associative reach of unaided thought with computational refinement.

4.3 Regression Modeling: Predictors of Originality

To identify the internal behavioral mechanisms driving originality outcomes, a hierarchical multiple regression was conducted on participants utilizing AI ($n = 320$, pooling Conditions 2 and 3).

Table 3: Hierarchical Regression Predicting Conceptual Originality

Predictor Variables	Model 1 (β)	Model 2 (β)	Model 3 (β)
Prior GPA	.112*	.084	.062
Digital / AI Literacy Score	.145**	.105*	.071
Unconstrained Prompt Count	---	-.382*	-.241
Iterative Counter-Prompts (%)	---	.415*	.389*
Verbatim AI Text Retention (%)	---	---	-.442***
Pre-AI Ideation Time (Minutes)	---	---	.318***
R^2	.048	.312	.468
Adjusted R^2	.042	.303	.458
ΔR^2	.048*	.264*	.156*
F for ΔR^2	7.98	40.28*	45.92***

* p < .05, **p < .01,** * p < .001

The regression analysis indicates that the *Verbatim AI Text Retention Percentage* ($\beta = -.442, p < .001$) and *Unconstrained Prompt Count* ($\beta = -.241, p < .01$) are strong negative predictors of creative originality.

Conversely, *Iterative Counter-Prompts* ($\beta = .389, p < .001$) and *Pre-AI Ideation Time* ($\beta = .318, p < .001$) accounted for 45.8% of the explained variance in student originality ($R^2 = .468, F = 45.92, p < .001$).

5. Qualitative Findings and Prompt Log Analysis

Analysis of the 36 interview transcripts and hundreds of participant prompt logs revealed three core cognitive archetypes that explain the statistical outcomes:

5.1 The Trap of the Immediate Plausible: Design Fixation

Unconstrained participants overwhelmingly engaged in "one-shot" prompting strategies. Typically, students entered prompts such as: *"Give me 5 creative ways to solve plastic pollution in the Mekong Delta."*

The model reliably returned structurally pristine, standard solutions: bio-degradable agricultural mulch, local government awareness campaigns, plastic collection drones, and community recycling incentives.

Prompt trace analysis revealed that once these lists were presented, 88.4% of unconstrained students adopted one of the provided five options as the foundation for their final submission. In retrospective interviews, Participant VN-U088 noted:

"The AI's response was so clean, well-structured, and fast that I immediately abandoned my initial fragmented thought about fungal mycoremediation tied to local shrimp farms. The AI's idea about robotic filtration skimmers seemed far more practical, so I just spent my remaining time expanding on that."

This confirms the presence of acute design fixation: exposure to fluent, high-probability algorithmic text anchors the human cognitive process, causing early termination of divergent exploration.

5.2 Algorithmic Echoes and Cultural Homogenization

A qualitative cross-comparison of the unconstrained essays revealed recurrent lexical choices and identical conceptual architectures. Themes such as "IoT-enabled monitoring stations," "Public-Private Partnerships (PPP)," and "Eco-credit mobile apps" appeared across 71% of essays in the unconstrained cohort.

Because LLMs prioritize standard global development vocabulary, the distinct socio-cultural peculiarities of rural Vietnamese riverine farming communities were systematically replaced by generic institutional rhetoric.

Participant VN-U214 reflected:

"When I wrote essays myself in previous semesters, my ideas were messy and drew on stories my grandparents told me about floods in An Giang. When I ask ChatGPT, it produces ideas that sound like an official UN report. It sounds smart, but it doesn't sound like anyone from the Delta actually thought of it."

5.3 Metacognitive Dialectic: Prompting as Socratic Conflict

In sharp contrast, students in the Structured Scaffolding condition used the AI not as an answering engine, but as an adversary. Their prompt logs revealed conversational friction:

- *"Here is my proposal to use indigenous aquatic ferns for microplastic absorption. Tear this proposal apart. Give me 4 technical reasons why this would fail in brackish water."*
- *"Do not give me generic solutions like recycling apps. What are three counter-intuitive, radical intersections between blockchain aquaculture verification and peasant credit co-ops?"*

These students used the machine to systematically stress-test edge cases, maintaining conceptual ownership while using algorithmic speed to expose logical vulnerabilities.

6. Discussion

The empirical findings from this investigation resolve several contentious debates regarding GenAI's place in higher education:

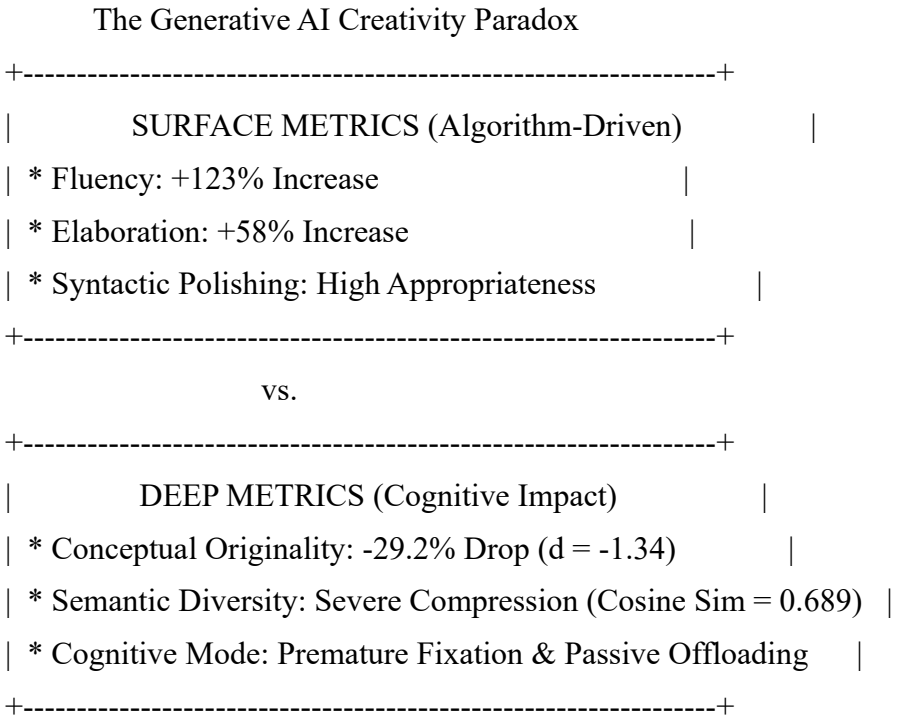

6.1 Deconstructing the Fluency-Originality Paradox

Our quantitative data demonstrates that unconstrained GenAI deployment induces a clear trade-off: it massively inflates task fluency while compressing conceptual originality. This empirical reality challenges early claims that unguided exposure to LLMs naturally elevates human creative output (Dwivedi et al., 2023).

Unconstrained access encourages cognitive path-of-least-resistance behavior. Generating truly novel insights requires tolerating cognitive ambiguity, dead-ends, and the mental labor of connecting distantly related concepts (Runco & Jaeger, 2012). When an LLM instantly presents a cohesive, highly articulate resolution, the intrinsic motivation to endure cognitive friction collapses. The student accepts an algorithmic local optimum, missing the creative breakthroughs located in the conceptual periphery.

6.2 The Cognitive Anchor-Rebound Model (CARM)

To conceptualize how instructional scaffolding successfully counteracts algorithmic fixation, we propose the *Cognitive Anchor-Rebound Model* (CARM).

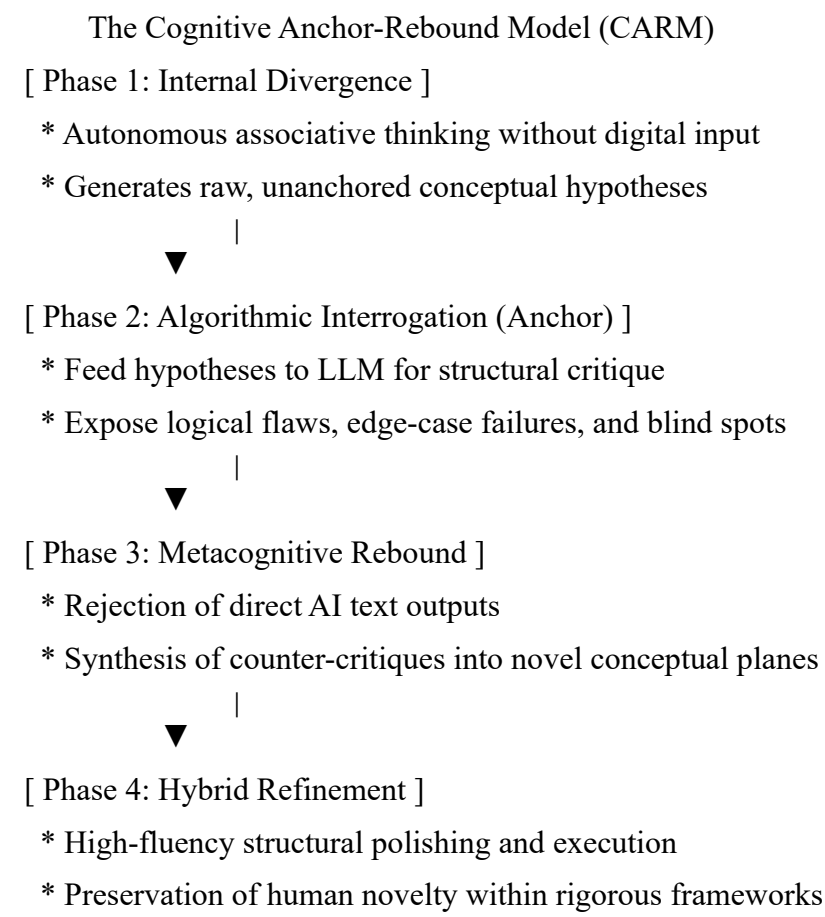

Under CARM, the AI is explicitly blocked from acting as the ideational originator. By enforcing Phase 1 (pre-AI ideation), students establish an internal cognitive anchor rooted in their idiosyncratic knowledge base and lived experiences. When the AI is introduced in Phase 2, its role is restricted to that of an analytical foil.

The student then "rebounds" off the algorithmic critique, achieving conceptual novelty that neither an unaided student (constrained by working memory limits) nor an unguided AI user (trapped in statistical conformity) could attain.

6.3 Implications for the Vietnamese Pedagogical Ecosystem

In Vietnam's evolving higher education landscape, these findings are critical. The current institutional reflex ranges from unmonitored integration to ineffective outright bans.

Neither extreme serves students. Bans ignore the reality of ubiquitous mobile access and international workplace demands, while unmonitored adoption risks entrenching intellectual passivity within a cultural context that is still transitioning away from rote knowledge replication.

Vietnamese universities must integrate AI literacy not as a study in syntax or interface management, but as a discipline of metacognitive self-regulation and argumentative inquiry.

7. Actionable Institutional and Classroom Frameworks

To operationalize these findings for university educators, we present the **C-A-R-E Instructional Framework**, tailored for integration across continuous assessment, studio critiques, and seminar modules:

Table 4: The C-A-R-E Pedagogical Implementation Protocol

Stage	Instructional Mandate	Student Deliverable
C - Conceptualize (Unaided)	No-tech ideation. Students map raw associative networks, personal hypotheses, and edge ideas.	Handwritten or offline concept sketches and research questions prior to accessing any LLM.
A - Adversarial Interrogation	Students prompt the AI as a skeptical referee. Direct prompts for answers are banned; prompt must seek critique.	Verbatim prompt-log showing interrogation of assumptions, counter-examples, and blind spots.
R - Rebound & Diverge	Students systematically identify what the AI got wrong, missed, or stated too generically, then diverge outward.	Comparative synthesis memo: "What the machine missed and how my final concept diverges."
E - Execute & Refine	Iterative deployment of LLM for syntax, structural organization, and code polishing; student owns the thesis.	Final submitted paper/artifact with an explicit Appendix of Algorithmic Contribution.

Assessment Modernization Guidelines

Institutions must shift assessment weights away from the final, polished written artifact—which an LLM can generate effortlessly—toward the visible provenance of thought:

1. **Mandatory Submission of Prompt Provenance:** Grade the intellectual rigor, depth, and adversarial sophistication of the prompt dialogue rather than the generated prose alone.

2. **Oral Defenses and Conceptual Inoculation:** Pair written submissions with rapid, unassisted five-minute oral cross-examinations to verify whether the student possesses deep conceptual ownership of the creative leaps within their project.

8. Limitations and Directions for Future Research

Several limitations should be noted:

1. **Temporal Scope:** This study evaluated immediate creative output within laboratory conditions. Longitudinal studies are needed to evaluate whether extended reliance on GenAI across multi-year degree programs leads to cumulative cognitive atrophy or, conversely, enhanced cognitive endurance.

2. **Model Evolution:** Experiments were conducted using frontier models available in 2024–2026 (GPT-4o, Claude 3.5 Sonnet). As reasoning architectures (e.g., OpenAI o1/o3-series) advance, their capacity to simulate divergent associative leaps may alter the observed fixation dynamics.

3. **Disciplinary Granularity:** While our sample spanned three academic tracks, finer-grained disciplinary breakdowns (e.g., pure mathematics vs. speculative design) are warranted to establish domain-specific thresholds for cognitive scaffolding.

9. Conclusion

Generative AI is neither an unqualified creative savior nor an inevitable driver of intellectual decline; its pedagogical trajectory depends on the structure of cognitive engagement. When

universities allow unconstrained, passive adoption, GenAI functions as a cognitive anchor that compresses originality, homogenizes student expression, and induces design fixation under a veneer of computational fluency.

When wrapped in structured metacognitive scaffolding, however, GenAI operates as an effective cognitive foil. By encouraging students to establish unassisted hypotheses before engaging the system, and by training them to interrogate algorithmic outputs through adversarial dialogue, educators can cultivate students who leverage algorithmic scale while preserving the originality and contextual authenticity of human thought. For higher education in Vietnam and globally, the pedagogical imperative is clear: we must stop using AI to bypass intellectual effort, and begin using it to elevate the rigor of creative inquiry.

References

Amabile, T. M. (1996). *Creativity in context: Update to the social psychology of creativity*. Westview Press.

Baidoo-Anu, D., & Ansah, L. O. (2023). Education in the era of Generative Artificial Intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI*, 7(1), 52–62. https://doi.org/10.61969/jai.1337146

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM FAccT Conference*, 610–623. https://doi.org/10.1145/3442188.3445922

Cotton, D. R., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2), 228–239. https://doi.org/10.1080/14703297.2023.2190148

Do, T. H., Nguyen, H. M., & Le, D. T. (2024). Digital transformation in Vietnamese higher education: Strategic policy directives and pedagogical realities. *Higher Education Policy and Development in Southeast Asia*, 19(2), 114–132.

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose#