

Generative AI and Cognitive Mediation in Higher Education: A Computational Literature Review

Saeed Ahmad

University of Eastern Finland

Muhammad Amir Saeed

m.saeed@campus.unimib.it

University of Milano-Bicocca

Research Article

Keywords: Generative artificial intelligence, Higher education, Cognitive mediation, Critical thinking, Metacognition, Cognitive offloading

Posted Date: October 22nd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11427530/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Generative artificial intelligence (GenAI) is increasingly integrated into higher education, raising questions about how cognitively substantive activity is distributed between learners and AI systems. Yet research on these cognitive implications remains fragmented across constructs and disciplinary conversations. This study maps the cognitive landscape of GenAI research in higher education through a computational literature review of 1,435 Scopus-indexed publications from 2023 to 2026. The analysis combined bibliometric mapping, theory-informed dictionary-based detection of ten cognitive constructs, temporal prevalence analysis, document-level co-occurrence and phi association analysis, and Structural Topic Modelling (STM). Critical thinking was the most prominently represented construct, followed by metacognition, while automation bias, cognitive load, memory and attention, cognitive dependency, and cognitive offloading formed less prominent but differentiated concerns. Although critical thinking and metacognition showed the highest raw co-occurrence, automation bias and cognitive dependency exhibited the strongest positive prevalence-adjusted association, demonstrating that construct prominence and relational specificity represent different properties of the literature. STM identified eight thematic structures and indicated increasing relative prominence of cognitive and theoretical modelling over the study period. Integrating these patterns, the study proposes cognitive mediation as an interpretive framework encompassing cognitive engagement, regulation, redistribution, and vulnerability. The findings suggest that a central challenge for higher education is not simply whether GenAI supports or undermines learning, but how cognitive responsibility is distributed and regulated in learner–AI interaction. The study provides a basis for research and pedagogical design concerned with preserving meaningful learner agency while integrating GenAI into higher-education learning and assessment.

1. Introduction

Generative artificial intelligence (GenAI) has moved rapidly into higher education, supporting activities such as writing, feedback, information synthesis, problem solving, and research-related work. Its educational significance, however, extends beyond efficiency or technological adoption. Because GenAI can perform tasks that ordinarily require substantial learner-generated reasoning and judgement, its growing use has intensified questions about what cognitive work students continue to perform themselves, what they delegate to external systems, and how that distribution is regulated. Recent reviews accordingly frame GenAI as capable of supporting higher-order learning under structured conditions while also raising concerns about overreliance, cognitive offloading, weakened analytical autonomy, and reduced learner agency when use becomes poorly regulated or substitutive (Alubthane, 2026; Li et al., 2026; Wang et al., 2026).

This cognitive dimension is especially important in higher education because many of the activities GenAI can now support—argumentation, critical evaluation, synthesis, explanation, drafting, and problem solving—are also central to the development and assessment of advanced learning. Current evidence does not support a simple conclusion that GenAI either enhances or diminishes cognition. A systematic review of 67 empirical studies found that ChatGPT-supported critical and creative thinking was most

consistently associated with inquiry-oriented and scaffolded designs involving metacognitive regulation, argumentative reasoning, and reflective engagement, whereas unstructured use was more frequently associated with cognitive offloading and reduced epistemic engagement (Li et al., 2026). Similarly, a review of 89 university studies concluded that GenAI may function as a cognitive amplifier under structured pedagogical conditions but as a cognitive substitute when use is insufficiently guided (Alubthane, 2026). These findings shift the educational problem from whether GenAI should be used towards how learner–AI interaction affects the organisation and regulation of cognitive activity.

Research addressing this problem is expanding but remains conceptually dispersed. Dedicated reviews now examine higher-order cognition, critical and creative thinking, cognitive offloading, learner agency, and critical autonomy. For example, Wang et al. (2026) synthesised 123 studies and identified both agency-supportive patterns—associated with self-regulation, reflective judgement and feedback engagement—and agency-eroding patterns involving offloading, dependence, overreliance and weakened judgement. Suazo et al. (2026) similarly review GenAI through the tension between cognitive offloading and critical autonomy, while recent bibliometric and systematic work has concentrated specifically on GenAI and critical thinking in higher education. These studies provide increasingly detailed accounts of particular cognitive domains, but the resulting evidence remains distributed across partly separate research conversations.

This fragmentation leaves three questions insufficiently resolved. First, comparatively little is known about the relative prominence and temporal development of multiple cognitive constructs within the same GenAI–higher education corpus. Second, construct-specific reviews provide limited insight into how cognitive concerns are represented in relation to one another across the wider literature. Critical thinking, metacognition, cognitive load, automation bias, cognitive offloading, dependency, memory and attention, learning depth, deskilling, and cognitive debt may overlap conceptually, but their field-level patterns of joint representation remain underexplored. Third, construct-level analysis has rarely been connected systematically with the broader thematic organisation of GenAI research, making it difficult to determine how cognitive concerns are embedded within areas such as institutional integration, academic tool use, student engagement, writing and feedback, and professional education. The unresolved problem is therefore not an absence of research on GenAI and cognition, but insufficient integration of its construct-level, relational, temporal, and thematic dimensions.

The present study addresses this gap through a computational literature review of 1,435 Scopus-indexed publications on GenAI, cognition, and higher education published between 2023 and 2026. It combines bibliometric and keyword analysis, theory-informed dictionary-based detection of ten cognitive constructs, document-level co-occurrence and phi association analysis, and Structural Topic Modelling (STM). These complementary methods examine how the literature has developed, which cognitive constructs are most prominently represented, how construct representation changes over time, how constructs are associated at the document level, and within which broader thematic structures they are situated. Because cognition-related terminology formed part of corpus retrieval, the study maps the

internal cognitive landscape of a deliberately cognition-oriented literature rather than estimating the prevalence of cognitive research across all GenAI scholarship in higher education.

The findings are interpreted through a cognitive mediation perspective, which treats GenAI as an external system capable of participating in the distribution and regulation of cognitively substantive activity. Cognitive mediation is used here as an organising perspective rather than as a validated psychological model, maintaining a distinction between patterns in scholarly discourse and actual learner-level cognitive effects. The framework is developed in Section 2.3.

The study therefore contributes an integrated account of how cognitive concerns surrounding GenAI are developing within higher-education scholarship and how those concerns may inform emerging questions for teaching, learning, assessment, and institutional practice. The study thereby provides an integrated account of the cognitive landscape surrounding GenAI in higher-education scholarship and its implications for teaching, learning, assessment, and institutional practice.

Accordingly, the study addresses four research questions:

RQ1. How has research on GenAI and cognition in higher education developed between 2023 and 2026, and what conceptual patterns characterise this literature?

RQ2. Which cognitive constructs are most prominently represented in the literature, and how has their relative prominence changed over time?

RQ3. What patterns of association emerge among cognitive constructs in the GenAI–higher education literature?

RQ4. What latent thematic structures characterise the literature, and how has their relative prevalence changed over time?

2. Literature and Conceptual Framework

2.1 GenAI and the cognitive conditions of higher-education learning

The educational significance of generative artificial intelligence (GenAI) lies partly in its capacity to participate in activities that have traditionally required substantial learner-generated cognitive effort. Unlike technologies used primarily for information retrieval or storage, contemporary GenAI systems can generate explanations, synthesise information, construct arguments, provide feedback, and contribute to problem solving and written production. This creates opportunities for cognitive support but also complicates the relationship between task performance and learning. Recent reviews consequently suggest that the cognitive implications of GenAI depend less on use per se than on how technological assistance is incorporated into learning activity. Scaffolded and inquiry-oriented use may support

reasoning, reflection, and higher-order thinking, whereas poorly regulated use may encourage cognitive offloading, overreliance, and reduced epistemic engagement (Li et al., 2026; Wang et al., 2026).

Critical thinking and metacognition are particularly important within this emerging literature. Critical thinking concerns learners' capacity to evaluate information, evidence, arguments, and conclusions rather than accepting generated output uncritically. Recent reviews indicate that GenAI can support critical thinking when learners remain actively involved in questioning, comparison, verification, and reflective evaluation, but that unstructured reliance may reduce the cognitive effort involved in these processes (Cong-Lem & Nguyen, 2026; Li et al., 2026). Metacognition addresses a related but distinguishable problem: whether learners continue to monitor, evaluate, and regulate their own cognitive activity when an external system participates in the task. Recent learning-to-learn research therefore positions metacognitive and self-regulatory capacities as important safeguards against GenAI overreliance and as conditions for maintaining learner agency in AI-supported environments (Schorr et al., 2026; Wang et al., 2026).

2.2 Cognitive redistribution, offloading, and vulnerability

These concerns can be situated within a broader understanding of cognition as activity that may be distributed across individuals and external resources. Sociocultural theory conceptualises higher psychological functions as mediated through culturally constituted tools and signs (Vygotsky, 1978), while distributed-cognition and extended-mind perspectives challenge the assumption that cognitively relevant activity occurs exclusively within the individual (Hutchins, 1995; Clark & Chalmers, 1998). Cognitive offloading provides a more specific account of this process by describing the use of external action or resources to reduce the internal cognitive demands of a task (Risko & Gilbert, 2016). These perspectives do not imply that technological assistance is inherently detrimental. Rather, they direct attention towards how cognitive work is distributed between learner and external system.

This distinction becomes especially consequential with GenAI because external support may range from scaffolding learner-generated thought to substituting for processes through which learning is intended to occur. A recent scoping review of 123 studies found a dual pattern: GenAI use may support learner agency through self-regulation, reflective judgement, feedback engagement, and related processes, while poorly regulated use is associated in the literature with offloading, overreliance, dependence, uncritical uptake, and weakened judgement (Wang et al., 2026). Related work distinguishes between forms of offloading that preserve cognitive agency and those in which core cognitive activity is increasingly delegated to AI (Wang et al., 2026; Suazo et al., 2026). The educational issue is therefore not simply whether learners offload cognitive work, but what they offload, for what purpose, and whether they retain responsibility for monitoring and evaluating the resulting output.

Automation bias and cognitive dependency extend this problem from redistribution towards vulnerability. Automation bias refers to inappropriate reliance on automated recommendations or outputs and has long been studied in human–automation interaction (Parasuraman & Riley, 1997). In GenAI-supported learning, the concern is that fluent and apparently authoritative generated responses may encourage

learners to defer evaluation or judgement to the system. Dependency represents a broader possibility in which repeated external assistance becomes increasingly substitutive rather than supportive. Emerging scholarship similarly connects cognitive offloading with questions of epistemic trust and learning accountability, reinforcing the importance of examining who retains cognitive and epistemic responsibility in AI-mediated learning (Suazo et al., 2026; Wang et al., 2026). These patterns should not be treated as evidence that GenAI necessarily causes dependency or cognitive decline; current reviews emphasise heterogeneous evidence and the importance of pedagogical conditions (Alubthane, 2026; Wang et al., 2026).

Longer-term concerns such as deskilling and cognitive debt extend this question further. They raise the possibility that repeated substitution of learner-generated cognitive activity could affect not only immediate task engagement but also the maintenance or development of capabilities over time (Yang et al., 2026). At present, however, evidence for such longer-term consequences remains considerably less established than evidence concerning critical thinking, metacognition, offloading, and overreliance (Alubthane, 2026; Li et al., 2026; Wang et al., 2026). They are therefore better treated as emerging research concerns requiring longitudinal and process-based investigation, rather than established effects of GenAI use.

2.3 Cognitive mediation as an integrative perspective

Taken together, these strands suggest that the central cognitive question posed by GenAI is neither simple enhancement nor simple impairment. It concerns the mediation and distribution of cognitive activity when learners interact with systems capable of performing cognitively substantive parts of academic tasks. The present study uses cognitive mediation as an integrative perspective for organising this problem. The concept draws broadly on mediation and distributed-cognition traditions but is used here as an interpretive framework, not as a validated psychological model or causal theory.

Four analytically related dimensions organise the framework. Cognitive engagement concerns learners' active involvement in reasoning, evaluation, attention, memory, and depth of processing, with critical thinking representing a prominent expression of this dimension. Cognitive regulation concerns the monitoring and control of learning activity, particularly through metacognitive and self-regulatory processes. Cognitive redistribution concerns the allocation of cognitive work between learner and technological system and encompasses cognitive load and cognitive offloading. Cognitive vulnerability concerns circumstances in which redistribution may become insufficiently regulated or increasingly substitutive, bringing automation bias, cognitive dependency, deskilling, and cognitive debt into view.

These dimensions are conceptually connected but should not be assumed to constitute empirically discrete clusters. A construct may contribute to more than one dimension, and theoretical proximity does not require strong statistical association within the literature. The framework therefore does not predetermine the co-occurrence or topic-modelling results. Instead, it provides an interpretive structure through which patterns identified computationally can be related to the broader educational question of how cognitive responsibility is distributed and regulated in learner–GenAI interaction. This distinction is

important because the present study analyses representations of cognition in scholarly publications rather than measuring cognitive processes in learners directly. The framework consequently supports interpretation of the research landscape while leaving learner-level causal relationships for empirical investigation.

3. Methodology

3.1 Research design

This study employed a computational literature review to examine how cognitive constructs are represented, associated, and thematically organised within research on generative artificial intelligence (GenAI) in higher education. The unit of analysis was the scholarly literature rather than individual learners or educational interventions. Accordingly, the study maps patterns in research discourse and does not infer cognitive effects of GenAI use.

The analytical design combined four complementary components: (1) descriptive bibliometric and keyword analysis to characterise corpus development and conceptual patterns; (2) theory-informed dictionary-based detection to examine the representation and temporal prominence of selected cognitive constructs; (3) document-level co-occurrence and phi association analysis to investigate relationships among those constructs; and (4) Structural Topic Modelling (STM) to identify broader latent thematic structures and their temporal development. Bibliometric analyses were conducted in R using the bibliometrix package (Aria & Cuccurullo, 2017), while topic modelling was conducted using the stm package (Roberts et al., 2014).

3.2 Data source, search strategy and corpus construction

Bibliographic records were retrieved from Scopus because its multidisciplinary coverage encompasses education, psychology, computer science, technology, health sciences, and other fields relevant to GenAI and higher education and provides structured metadata suitable for bibliometric and computational analysis.

The search strategy was developed iteratively around three conceptual domains: GenAI, higher education, and cognition and learning. GenAI terms included expressions referring to generative artificial intelligence, ChatGPT, large language models, and AI assistants; higher-education terms captured university, tertiary, and related contexts; and the cognition-and-learning component combined broader cognitive terminology with constructs relevant to AI-mediated learning, including critical thinking, metacognition, cognitive load, cognitive offloading, automation bias, cognitive dependency, deskilling, memory, attention, and cognitive debt. The complete Scopus search string is provided in the supplementary material.

The search covered 2022–2026 and was conducted on 27 June 2026. No eligible records were retrieved for 2022; the analytical corpus therefore spans 2023–2026. Because 2026 represented only publications

indexed by the search date, it was treated as a partial year in all temporal interpretations. Records were eligible when they addressed GenAI in a higher-education context, fell within the specified period, were indexed in Scopus, and were published in English. The retrieved scholarly document types were retained because the purpose was computational mapping of the research landscape rather than synthesis of intervention effects.

The search returned 1,503 records. During bibliometrix conversion, 67 duplicate records were removed. DOI and normalised-title checking identified one further duplicate, resulting in a final bibliometric corpus of 1,435 unique publications. One record containing insufficient textual material was subsequently excluded from STM preprocessing, producing an STM corpus of 1,434 documents. These corpus-construction procedures correspond to the cleaning process documented in the full analytical workflow.

3.3 Text preparation and cognitive-construct analysis

Titles, abstracts, author keywords, and index keywords were combined at document level to create the textual corpus. General preprocessing included conversion to lower case, tokenisation, punctuation removal, and stop-word removal. Stemming or lemmatisation was not applied to the general text-processing pipeline. STM-specific preprocessing, however, used `stm::textProcessor` with `stem = TRUE`. After removal of low-frequency terms, the STM corpus contained 1,434 documents, 2,265 retained terms, and 149,170 tokens.

A theory-informed lexical dictionary was developed to operationalise explicit textual references to ten cognitive constructs: critical thinking, metacognition, automation bias, cognitive load, memory and attention, cognitive dependency, cognitive offloading, learning depth, deskilling, and cognitive debt. Each construct was represented through manually specified regular-expression patterns incorporating relevant synonyms, spelling and hyphenation variants, morphological forms, and multi-word expressions. The complete dictionary and regular-expression patterns are provided in the supplementary material.

The dictionary was applied at document level. A construct was coded as present when at least one qualifying lexical match occurred within a document, producing binary presence/absence indicators for each construct. Construct prevalence was subsequently calculated as the number and proportion of documents in which each construct was detected and was examined by publication year to assess changes in relative representation.

The dictionary was designed as a transparent lexical operationalisation rather than a validated measurement instrument. No formal inter-rater or dictionary-validation procedure and no systematic manual exclusion of ambiguous matches were undertaken. Consequently, the analysis captures explicit lexical representation of constructs and should not be interpreted as an exhaustive measurement of their conceptual presence. Because cognition- and learning-related terminology also formed part of the original retrieval strategy, comparisons describe the relative prominence of constructs within a deliberately cognition-oriented corpus, rather than the prevalence of cognitive research across all GenAI scholarship in higher education. These boundaries are integral to the original analytical design.

3.4 Co-occurrence and association analysis

A binary document-by-construct matrix derived from the dictionary analysis was used to examine relationships among cognitive constructs. Two complementary measures were calculated. Raw co-occurrence frequencies indicated how often pairs of constructs appeared within the same publication, whereas phi coefficients measured the strength and direction of association between pairs of binary construct indicators while accounting for their presence/absence distributions.

Network visualisation was used to represent the broader structure of joint construct representation. Construct pairs with non-zero raw co-occurrence were retained, edge weights represented document-level co-occurrence frequencies, and the network was displayed using a Fruchterman–Reingold layout.

Raw co-occurrence and phi were interpreted separately because frequently occurring constructs can produce high joint frequencies without exhibiting strong positive prevalence-adjusted associations. Neither measure was interpreted as evidence of causal or learner-level cognitive relationships. This distinction is particularly important for interpreting the contrast between raw co-occurrence and adjusted association in the subsequent results.

3.5 Structural Topic Modelling and analytical integration

Structural Topic Modelling was used to identify latent thematic structures within the textual corpus. STM models documents as probabilistic mixtures of topics and topics as distributions over terms, enabling recurrent patterns of vocabulary to reveal broader thematic organisation (Roberts et al., 2014).

An eight-topic solution ($K = 8$) was specified for the 1,434-document STM corpus. The analytical workflow did not include formal optimisation of K through comparison of alternative solutions using semantic coherence, exclusivity, residuals, or held-out likelihood. The eight-topic solution is therefore treated as a researcher-specified modelling choice providing substantively interpretable thematic resolution, rather than as an empirically optimal number of topics. Topics were interpreted using highest-probability, FREX, lift, and score terms generated through `stm::labelTopics`, together with topic-term outputs. Representative high-probability documents were not systematically examined during labelling; topic labels should consequently be understood as substantive interpretations of probabilistically estimated lexical structures rather than predetermined or definitive thematic categories.

Publication year was included in the STM as a smooth topic-prevalence covariate ($\text{prevalence} = \sim s(\text{Year})$). Document-level topic proportions were also aggregated by year to examine changes in relative thematic prevalence. Trends involving 2026 were interpreted cautiously because the final year was incomplete at the time of retrieval.

The four analytical components were integrated only after their respective patterns had been established. Bibliometric analysis characterised corpus development and broad conceptual organisation; dictionary analysis mapped theoretically specified cognitive constructs; co-occurrence and phi analysis examined their document-level relationships; and STM identified broader latent thematic structures.

These outputs were subsequently interpreted through the cognitive mediation perspective developed in Section 2. The framework was therefore interpretive rather than confirmatory and was not used to infer learner-level cognitive effects from textual patterns.

4. Empirical Results

4.1 Dataset Description and Publication Development

The Scopus search yielded 1,503 records related to GenAI, higher education, and cognitive or learning-related constructs. Following duplicate detection during bibliometrix conversion, 67 records were removed. One additional duplicate was identified through DOI and normalised-title checking, resulting in a final bibliometric corpus of 1,435 unique publications. Although the original search window covered 2022–2026, no eligible records were retrieved for 2022; the analytical corpus therefore covered publications from 2023 to 2026. Following text preprocessing, one additional record containing insufficient textual material was excluded from the STM analysis. The final STM corpus comprised 1,434 documents, 2,265 retained terms, and 149,170 tokens.

Publication output increased substantially between 2023 and 2025. The corpus contained 44 publications in 2023, 225 in 2024, and 657 in 2025, representing increases of approximately 411% from 2023 to 2024 and 192% from 2024 to 2025. A further 509 publications were retrieved for 2026. Because the search was conducted on 27 June 2026, however, the 2026 count represents only a partial publication year and should not be interpreted as evidence of declining publication activity.

4.2 Conceptual Structure of the Literature

Keyword analysis indicated a broad conceptual distribution across technological, educational, academic, and cognitive concerns. Frequently occurring author and index keywords included AI, students, education, learning, critical, generative, thinking, artificial intelligence, ChatGPT, higher, academic, tools, skills, writing, language, cognitive, teaching, assessment, feedback, literacy, integration, engineering, model, engagement, and support.

The keyword co-occurrence network showed interconnected conceptual configurations around GenAI use in higher education. Prominent patterns linked ChatGPT, students, academic use, writing, and assessment; learning, teaching, educational design, and student engagement; and cognition, critical thinking, metacognition, cognitive load, and automation-related concerns. These patterns indicate that cognitive discourse is embedded within a wider research landscape concerned with both educational practice and GenAI-supported academic activity. The 30 most frequent index keywords are presented in Supplementary Figure S1.

4.3 Cognitive Construct Representation and Temporal Trends

Dictionary-based analysis revealed substantial differences in the representation of the ten cognitive constructs across the corpus. Critical thinking was the most frequently detected construct, occurring in 647 publications (45.1%), followed by metacognition (256; 17.8%), automation bias (212; 14.8%), cognitive load (198; 13.8%), memory and attention (184; 12.8%), cognitive dependency (176; 12.3%), and cognitive offloading (162; 11.3%). Learning depth (96; 6.7%), deskilling (63; 4.4%), and cognitive debt (41; 2.9%) were less frequently represented. The complete frequency distribution is shown in Supplementary Figure S2.

Temporal analysis showed that these constructs did not develop uniformly between 2023 and 2026. Critical thinking remained the most prominent construct throughout the period, while metacognition became increasingly visible as the literature expanded. Cognitive load and memory/attention also maintained a substantial presence across the corpus. Automation bias, cognitive dependency, and cognitive offloading were represented more frequently in absolute terms in the later years of the publication period.

The least frequently represented constructs—learning depth, deskilling, and cognitive debt—remained comparatively peripheral across the study period. Because the 2026 corpus includes only publications indexed up to 27 June, changes involving that year should be interpreted as partial-year patterns rather than complete annual trends.

4.4 Cognitive Construct Co-occurrence and Associations

Document-level co-occurrence analysis showed that several cognitive constructs were frequently represented within the same publications. The most frequent co-occurrence involved critical thinking and metacognition ($n = 231$), followed by critical thinking and automation bias ($n = 170$), critical thinking and cognitive dependency ($n = 88$), and critical thinking and memory and attention ($n = 83$). Automation bias and cognitive dependency ($n = 56$) also co-occurred relatively frequently.

Because raw co-occurrence frequencies are influenced by the overall prevalence of individual constructs, phi coefficients were used to examine the relative strength and direction of binary document-level associations. After accounting for construct prevalence, the strongest positive association was between automation bias and cognitive dependency ($\phi = .310$). Other positive associations were substantially weaker, including cognitive offloading–cognitive debt ($\phi = .138$), cognitive offloading–memory and attention ($\phi = .092$), cognitive load–memory and attention ($\phi = .086$), and critical thinking–automation bias ($\phi = .085$).

The strongest negative associations occurred between critical thinking and cognitive load ($\phi = -.460$) and critical thinking and metacognition ($\phi = -.342$). Weaker negative associations were observed between cognitive offloading and critical thinking ($\phi = -.144$) and between critical thinking and deskilling ($\phi = -.104$). These results show that frequent co-occurrence did not necessarily correspond to positive prevalence-adjusted association. Most notably, critical thinking and metacognition had the highest raw co-occurrence count ($n = 231$) but a moderately negative phi coefficient ($\phi = -.342$), whereas critical

thinking and automation bias had the second-highest raw co-occurrence (n = 170) but only a weak positive phi association ($\phi = .085$).

These associations represent patterns of joint representation within publications and should not be interpreted as relationships among learners’ cognitive processes.

4.5 Structural Topic Modelling

Structural Topic Modelling identified eight latent topics within the 1,434-document corpus. Topic interpretation was based on the highest-probability, FREX, lift, and score terms generated by the model. The resulting topics represented distinct but partly overlapping areas of GenAI research in higher education, encompassing higher education integration, academic tool use, cognitive-theoretical modelling, critical thinking, writing and feedback, learning engagement, professional education, and engineering and design applications.

Table 1
Interpretive labels for the eight STM topics

Topic	Interpretive label	Indicative focus
Topic 1	Higher education integration	GenAI adoption, institutional integration, higher education systems
Topic 2	Cognitive and theoretical modelling	Cognition, theory, design, models, cognitive load
Topic 3	Professional and health education applications	Artificial intelligence, nursing, medical and human-centred education
Topic 4	ChatGPT and critical thinking	ChatGPT, students, education, thinking, criticality, skills
Topic 5	Writing, language and feedback	Writing, language learning, feedback, English/EFL contexts
Topic 6	Student learning and engagement	Learning, interaction, engagement, groups and thinking
Topic 7	Academic tool use	Students, academic use, educational tools, GenAI practices
Topic 8	Engineering, design, and computational learning	Engineering, design, computing, models and learning

Higher education integration (Topic 1) had the highest overall prevalence. Other prominent topics included Academic tool use (Topic 7), Engineering, design and computational learning (Topic 8), ChatGPT and critical thinking (Topic 4), and Student learning and engagement (Topic 6). The remaining topics were Cognitive and theoretical modelling (Topic 2), Writing, language and feedback (Topic 5), and Professional and health education applications (Topic 3).

Temporal estimates indicated that topic prevalence varied across the 2023–2026 publication period. Because the corpus includes only publications indexed up to 27 June 2026, estimates for 2026 should be interpreted cautiously.

Cognitive and theoretical modelling (Topic 2) increased substantially in relative prevalence across the study period. Writing, language and feedback (Topic 5) and Student learning and engagement (Topic 6) also showed overall increases. In contrast, ChatGPT and critical thinking (Topic 4) and Engineering, design and computational learning (Topic 8) declined in relative prevalence across the period.

Higher education integration (Topic 1) declined after 2023 before increasing moderately in the partial 2026 corpus. Academic tool use (Topic 7) increased through 2025 before declining in the partial 2026 dataset. Professional and health education applications (Topic 3) declined initially before increasing again across 2025 and the partial 2026 corpus. Changes involving the final year should be interpreted descriptively because 2026 represents an incomplete publication year.

5. Discussion

This study examined how cognitive concerns are represented, associated, and thematically organised within the rapidly expanding literature on GenAI in higher education. Three findings are particularly important. First, the literature has developed from a relatively small publication base into a differentiated research landscape in which critical thinking remains highly prominent while metacognition, cognitive load, offloading, dependency, and automation-related concerns have become increasingly visible. Second, construct prominence and construct association produced different patterns: frequently represented constructs did not necessarily exhibit strong positive prevalence-adjusted relationships. Third, the topic structure indicates that cognitive concerns are embedded within broader questions of GenAI integration, academic practices, writing and feedback, student engagement, and disciplinary applications, while cognitive-theoretical modelling has gained relative prominence over time. Together, these findings suggest an increasing differentiation of research beyond questions of GenAI adoption towards how cognitive activity is distributed and regulated in AI-supported higher education.

5.1 From GenAI adoption to a diversified cognitive agenda

The rapid expansion of the corpus between 2023 and 2025 reflects the extraordinary speed with which GenAI became an object of higher-education research following the widespread availability of conversational large language models. More important than publication growth itself, however, is the diversification of the cognitive questions being investigated. Critical thinking was by far the most frequently represented construct, occurring in 45.1% of publications, followed by metacognition, automation bias, cognitive load, memory and attention, cognitive dependency, and cognitive offloading. This pattern is consistent with recent reviews positioning critical evaluation, reflective judgement, and metacognitive regulation as central to productive learner engagement with GenAI (Li et al., 2026; Wang et al., 2026).

The prominence of critical thinking is understandable. GenAI systems can produce fluent and plausible responses without guaranteeing their accuracy or epistemic reliability, making evaluation, verification, comparison, and judgement increasingly consequential aspects of academic AI use. Recent scholarship therefore positions critical thinking not simply as an educational outcome potentially affected by GenAI, but also as a capacity required for evaluating, verifying, and using AI-generated information responsibly (Cong-Lem & Nguyen, 2026; Li et al., 2026). Metacognition adds a regulatory dimension to this problem because productive AI-supported learning requires learners to monitor not only their own understanding but also when, why, and to what extent external assistance should be used (Li et al., 2026; Wang et al., 2026).

At the same time, the presence of automation bias, cognitive dependency, and cognitive offloading indicates that the literature increasingly addresses the possible consequences of transferring cognitive work to GenAI. Cognitive offloading is not inherently detrimental; external resources have long formed part of human cognitive activity (Risko & Gilbert, 2016). The educational issue is instead whether offloading reduces unnecessary cognitive demands while preserving learner engagement, or substitutes for processes through which knowledge, judgement, and competence are expected to develop. This distinction is consistent with recent evidence portraying GenAI as potentially supportive of learner agency under conditions of active monitoring and self-regulation, but potentially agency-eroding when reliance becomes substitutive or poorly regulated (Wang et al., 2026).

The comparatively low representation of learning depth, deskilling, and cognitive debt is also informative. The comparatively low representation of learning depth, deskilling, and cognitive debt is also informative. Their peripheral position within the corpus suggests that possible longer-term consequences of repeated cognitive delegation remain less developed in the literature than immediate concerns such as critical thinking, metacognition, and cognitive load.

5.2 Cognitive prominence is not the same as cognitive association

The association analysis provides a more differentiated picture than construct frequencies alone. Critical thinking and metacognition showed the highest raw co-occurrence ($n = 231$), but their prevalence-adjusted phi coefficient was moderately negative ($\phi = -.342$). Critical thinking and cognitive load showed an even stronger negative association ($\phi = -.460$). Conversely, automation bias and cognitive dependency exhibited the strongest positive phi association ($\phi = .310$), despite a considerably lower raw co-occurrence count.

This distinction matters conceptually as well as statistically. High raw co-occurrence can arise because two constructs are individually common within a corpus; it does not necessarily mean that they form a disproportionately associated research pairing. The negative phi coefficient for critical thinking and metacognition therefore should not be interpreted as evidence that the cognitive processes themselves are negatively related. Rather, within this literature, their document-level representation overlaps less than would be expected from their individual prevalence. This pattern may indicate that two theoretically

complementary aspects of AI-supported learning—critical evaluation and metacognitive regulation—are being represented within partly differentiated research conversations.

The automation bias–cognitive dependency association presents a different pattern. Its positive phi coefficient indicates a comparatively strong tendency for these constructs to be represented jointly within publications. Conceptually, this pairing is plausible: automation bias concerns inappropriate deference to automated outputs, whereas cognitive dependency concerns increasing reliance on external technological support. Established human–automation research has long identified inappropriate reliance as a central problem in automated decision environments (Parasuraman & Riley, 1997), while emerging GenAI scholarship increasingly raises related concerns about overreliance and diminished evaluative agency (Wang et al., 2026). The present analysis does not establish that automation bias produces cognitive dependency, but it shows that these concerns exhibit a comparatively strong pattern of joint representation within the present corpus.

This finding has a methodological implication for computational reviews as well. Reliance on frequency or raw co-occurrence alone would have produced a substantially different account of the field. Combining prevalence with adjusted association distinguishes visibility from relational specificity, revealing less frequent constructs that may nevertheless form comparatively coherent conceptual pairings.

5.3 From tool use towards cognitive mediation

The STM results place these construct-level patterns within the broader thematic organisation of the literature. Higher education integration remained the most prevalent topic overall, while academic tool use, engineering and computational applications, critical thinking, and student learning and engagement were also prominent. This indicates that the corpus remains strongly oriented towards questions of implementation and educational application. At the same time, the temporal trajectories suggest a changing thematic balance. Cognitive and theoretical modelling increased substantially in relative prevalence, alongside growth in writing, language and feedback and student learning and engagement, whereas the relative prevalence of the more explicitly ChatGPT-and-critical-thinking topic declined.

The decline of a specific ChatGPT/critical-thinking topic should not be read as declining scholarly concern with critical thinking; dictionary analysis shows that critical thinking remained the most prominent cognitive construct throughout the corpus. One possible interpretation is that critical thinking has become distributed across a broader range of GenAI-related themes rather than remaining concentrated within a single identifiable topic. Similarly, the increasing relative prevalence of cognitive-theoretical modelling is consistent with greater conceptual differentiation within the corpus.

Viewed through the cognitive mediation perspective, these patterns can be interpreted across four interconnected dimensions: cognitive engagement in learner-generated reasoning and evaluation; cognitive regulation through monitoring and control of AI use; cognitive redistribution through changes in

cognitive load and offloading; and cognitive vulnerability associated with automation bias, dependency, deskilling, and cognitive debt.

These dimensions should not be interpreted as empirically validated categories or as a sequential causal model. Rather, they provide an organising framework for a literature in which GenAI is increasingly discussed not merely as an educational tool but as a participant in the distribution of cognitive work. This shifts the central higher-education question away from whether GenAI is intrinsically beneficial or harmful. The more consequential question is which forms of cognitive activity should be supported by GenAI, which may appropriately be redistributed, which should remain substantially learner-generated, and what forms of regulation are required to preserve meaningful learning and learner agency.

6. Implications and Future Research

6.1 Implications for Higher Education

Higher-education practice should make the purpose, extent, and consequences of cognitive delegation explicit in GenAI-supported learning. Tasks involving GenAI should require students to remain responsible for evaluating information, justifying decisions, monitoring their understanding, and determining when external assistance supports rather than substitutes for learning.

The prominence of critical thinking and metacognition has particular implications for curriculum and assessment. If GenAI can contribute to drafting, synthesis, explanation, and problem solving, assessment designs that evaluate only the resulting product may provide increasingly limited evidence of the cognitive processes through which that product was produced. Assessment can instead incorporate activities that make reasoning and regulation more visible, such as justification of AI use, comparison and verification of generated outputs, reflection on prompting and revision decisions, and explanation of what cognitive work was retained or delegated. This approach is consistent with recent evidence emphasising critical evaluation, metacognitive monitoring, and self-regulation as important conditions for productive GenAI-supported learning (Li et al., 2026; Wang et al., 2026).

The positive document-level association between automation bias and cognitive dependency further suggests that AI literacy should encompass more than technical competence or prompt construction. Students need opportunities to develop evaluative AI literacy, what might be termed evaluative AI literacy: the capacity to question generated responses, recognise uncertainty and possible error, compare AI output with other sources, and retain responsibility for academic judgement. For teachers, the corresponding challenge is to design GenAI-supported activities in which technological assistance functions as scaffolding rather than an automatic replacement for learner-generated reasoning. At institutional level, guidance on GenAI could similarly complement academic-integrity provisions with explicit attention to cognitive engagement, learner agency, and responsible delegation.

These implications should nevertheless be treated as pedagogical propositions derived from patterns in the literature, not as intervention effects demonstrated by the present study. The analysis maps scholarly

discourse and therefore cannot establish which particular teaching or assessment practices produce better cognitive outcomes.

6.2 Future Research

The findings identify several priorities for future research. First, the strong representation of critical thinking alongside metacognition, cognitive load, offloading, dependency, and automation bias points to the need for studies examining these constructs relationally rather than individually. The contrast between raw co-occurrence and phi associations is particularly instructive: constructs that appear frequently in the same literature are not necessarily strongly associated after prevalence is taken into account. Future empirical research should investigate whether theoretically related constructs such as critical thinking and metacognition interact in actual learner–GenAI activity and under what pedagogical conditions those relationships become educationally consequential.

Second, greater attention is needed to the temporal dynamics of cognitive delegation. Learning depth, deskilling, and cognitive debt were comparatively peripheral in the present corpus, yet they concern consequences that may become visible only through sustained GenAI use. Cross-sectional studies and short interventions provide limited evidence for determining whether repeated cognitive offloading changes learners' independent performance, strategy use, knowledge retention, or capacity to complete comparable tasks without AI support. Longitudinal and repeated-measures designs are therefore particularly important.

Third, future studies should examine boundary conditions separating productive cognitive support from cognitive substitution. Metacognitive regulation is especially relevant because it may help explain why similar forms of GenAI assistance produce different patterns of learner engagement. Experimental and process-oriented research could examine how task complexity, prior knowledge, AI literacy, disciplinary context, assessment conditions, and degree of learner control influence decisions to accept, reject, verify, modify, or delegate cognitive work.

Finally, the field requires stronger methodological integration. Self-report measures remain useful for perceptions and attitudes but cannot independently establish how cognition changes during learner–AI interaction. Future research could combine performance measures, process data, interaction logs, stimulated recall, think-aloud protocols, longitudinal assessment, and validated measures of cognitive and metacognitive constructs. Such work would enable the cognitive mediation perspective proposed here to be tested, refined, or challenged empirically, rather than remaining primarily an interpretive framework for organising the literature.

7. Limitations

Several limitations should be considered when interpreting the findings. First, the corpus was restricted to English-language publications indexed in Scopus. Although Scopus provides broad multidisciplinary coverage, relevant research indexed elsewhere or published in other languages may therefore be

underrepresented. The corpus also covers a rapidly developing field over a relatively short period (2023–2026), and the 2026 data include only publications indexed up to 27 June. Temporal patterns involving the final year should consequently be interpreted as provisional.

Second, the search strategy deliberately incorporated cognition- and learning-related terminology to construct a corpus relevant to the study objectives. The reported prevalence of cognitive constructs therefore represents their relative prominence within a cognition-oriented GenAI–higher education corpus, rather than the prevalence of cognitive concerns across the entire GenAI-in-higher-education literature.

Third, cognitive constructs were identified through a theory-informed lexical dictionary applied to bibliographic text. Although the dictionary incorporated synonyms, morphological variants, and multi-word expressions, lexical detection cannot fully capture contextual meaning, implicit conceptualisation, or terminological ambiguity. The dictionary was not subjected to formal validation or systematic manual coding against a reference sample. Construct frequencies and associations should therefore be interpreted as patterns of explicit textual representation, rather than direct measurements of the underlying cognitive phenomena.

Fourth, the eight-topic STM solution was researcher specified rather than selected through systematic comparison of alternative topic solutions, and topic labels were interpreted primarily from model-generated term distributions. The topics should consequently be regarded as analytically useful representations of thematic structure rather than definitive categories of the field.

Finally, the study analyses scholarly discourse rather than learner cognition itself. Co-occurrence, phi associations, construct prevalence, and topic trajectories cannot establish causal relationships among cognitive processes or determine whether GenAI use enhances or impairs learning. The cognitive mediation perspective developed from these patterns should therefore be understood as an interpretive framework requiring subsequent empirical testing and refinement.

8. Conclusion

This study mapped the developing cognitive landscape of GenAI research in higher education through an integrated computational analysis of 1,435 Scopus-indexed publications from 2023 to 2026. The findings show that cognitive concerns have become increasingly differentiated within this rapidly expanding literature. Critical thinking occupies a particularly prominent position, alongside metacognition, cognitive load, automation bias, cognitive dependency, cognitive offloading, and related concerns. At the same time, the contrast between raw co-occurrence and prevalence-adjusted associations demonstrates that construct visibility does not necessarily correspond to relational specificity, highlighting the value of examining cognitive concerns as an interconnected rather than merely frequency-based research landscape.

The thematic analysis further indicates that cognition is embedded within broader discussions of GenAI integration, academic practices, writing and feedback, student engagement, and disciplinary applications, while cognitive and theoretical modelling has gained relative prominence over the study period. Taken together, these patterns support the cognitive mediation perspective developed in this study as an interpretive structure for understanding the distribution and regulation of cognitive activity in learner–GenAI interaction. These dimensions are proposed as an interpretive structure rather than a validated causal model.

For higher education, the central issue is therefore not simply whether GenAI supports or undermines learning, but how cognitive responsibility is distributed between learners and AI systems, how that distribution is regulated, and which forms of cognitive activity should remain meaningfully learner-generated. Addressing these questions through longitudinal, process-oriented, and pedagogically situated research will be important as GenAI becomes increasingly integrated into higher-education learning and assessment.

Declarations

Author Contribution

Author Contributions Statement S.A. (Saeed Ahmad) contributed to the original manuscript drafting, review, and editing, methodology development, validation, and visualization. M.A.S. (Muhammad Amir Saeed) contributed to software development, computational analysis, methodology development, empirical evaluation, validation, and visualization. All authors reviewed and approved the final version of the manuscript and agreed to be accountable for its content.

References

1. Alubthane, F. O. (2026). Amplifier or substitute? A systematic review of generative AI's impact on higher-order cognitive skills among university students. *Frontiers in Psychology*, 17, 1863931. 10.3389/fpsyg.2026.1863931
2. Aria, M., & Cuccurullo, C. (2017). bibliometrix: An R-tool for comprehensive science mapping analysis. *Journal of Informetrics*, 11 No(4), 959–975. 10.1016/j.joi.2017.08.007
3. Clark, A., & Chalmers, D. (1998). The extended mind, *Analysis*, Vol. 58 No. 1, pp. 7–19. 10.1093/analys/58.1.7
4. Cong-Lem, N., & Nguyen, T. T. (2026). Towards a framework for understanding and assessing critical thinking in generative AI-enhanced language learning: A scoping review. *European Journal of Psychology of Education*. 10.1007/s10212-026-01089-y
5. Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, 34 No(10), 906–911. 10.1037/0003-066X.34.10.906
6. Hutchins, E. (1995). *Cognition in the Wild*. MIT Press. 10.7551/mitpress/1881.001.0001

7. Li, C., Cui, H., & Hagedorn, L. S. (2026). The cognitive impact of ChatGPT in higher education: A systematic review of critical and creative thinking outcomes. *Computers and Education: Artificial Intelligence*, 10, 100571. 10.1016/j.caeai.2026.100571
8. Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *Psychology of Learning and Motivation*, 26, 125–173. 10.1016/S0079-7421(08)60053-5
9. Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39 No(2), 230–253. 10.1518/001872097778543886
10. Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20 No(9), 676–688. 10.1016/j.tics.2016.07.002
11. Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B., & Rand, D. G. (2014). Structural topic models for open-ended survey responses. *American Journal of Political Science*, 58 No(4), 1064–1082. 10.1111/ajps.12103
12. Schorr, I., Bardach, L., Bühler, B., & Kasneci, E. (2026). Learning-to-learn in the age of generative AI: A scoping review and conceptual framework. *Computers and Education: Artificial Intelligence*, 10, 100575. 10.1016/j.caeai.2026.100575
13. Suazo, I., Santiesteban Velázquez, M., Saracostti, M., & Chaple Gil, A. M. (2026). Between cognitive offloading and critical autonomy: A systematic review of the epistemic implications of generative AI in higher education. *Frontiers in Education*, 11, 1878665. 10.3389/feduc.2026.1878665
14. Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Harvard University Press.
15. Wang, G., Wang, W., Yang, D., & Ren, J. (2026). Generative AI, cognitive offloading, and learner agency in higher education: A scoping review. *Behavioral Sciences*, 16 No(7), 1150. 10.3390/bs16071150

Figures




Figure 1

Annual Publication Trend




Figure 2

Keyword Co-occurrence Network




Figure 3

Cognitive Construct Heatmap by Year

Network of Cognitive Construct Co-occurrences




Figure 4

Network of Cognitive Construct Co-occurrences




Figure 5

Phi Association Matrix of Cognitive Constructs




Figure 6

