

An Analytical Scoping Review of How Artificial Intelligence Reliance Deprives Academics of Creative Synthesis

Parhlad Singh Ahluwalia, Dr. Bhim Rao Ambedkar Law University, Jaipur, Rajasthan

Abstract

Background: Generative artificial intelligence (AI) is now embedded in academic reading, writing, and ideation. Its benefits for speed and fluency are well publicised. Its costs to the higher-order intellectual act that defines scholarship, creative synthesis, are far less examined. Creative synthesis means combining distant ideas, tolerating uncertainty, and producing a new conceptual arrangement that neither the sources nor the scholar held beforehand.

Objective: This paper analyses how habitual reliance on AI systems can deprive academics of the conditions under which creative synthesis develops, and it appraises the strength of the evidence for this claim.

Methods: A PRISMA-informed scoping review was conducted in two strands. Strand A was a transparent, title-level screening of the 117-record corpus supplied for this study. Strand B was a purposive search for theoretical and empirical anchor studies in cognitive psychology, creativity research, and AI ethics. Records were charted against five mechanisms of deprivation.

Results: Twenty-nine corpus records and eighteen external sources were retained ($n = 47$). Evidence converged on five mechanisms: cognitive offloading, the loss of productive difficulty, automation complacency, output homogenisation, and epistemic narrowing. Experimental evidence shows that AI assistance raises individual output quality while reducing the diversity of ideas across a group, and that performance gains during assisted work do not always survive its removal.

Conclusions: The evidence does not show that AI use is inherently harmful. It shows that unstructured, substitutive reliance erodes the effortful integrative processes on which synthesis depends. The paper closes with an evidence-graded framework for scholars and institutions.

Keywords: artificial intelligence; creative synthesis; cognitive offloading; academic integrity; higher education; scoping review

1. Introduction

1.1 The problem: fluency without formation

Scholarship has always been more than the accumulation of information. Its distinctive act is synthesis, the deliberate joining of findings, concepts, and methods from separate domains into a coherent contribution that did not exist before. Boden's influential account distinguishes combinational, exploratory, and transformational creativity, and it treats the first as the

recombination of familiar ideas in unfamiliar ways [1]. Koestler described the same event as bisociation, the collision of two habitually unconnected frames of reference that yields a novel insight [2]. On both accounts, synthesis is not a clerical operation. It is a formative experience in which the scholar's own mind is changed by the effort of connecting things.

Large language models now offer to perform this operation on request. A researcher can ask for a literature summary, a research gap, a theoretical framework, or a discussion section, and receive a fluent, structured answer within seconds. The product looks like the product of synthesis. The concern this paper develops is that the *process* which once produced both the product and the scholar is being bypassed. When the process is bypassed repeatedly, the capacity it builds may weaken. This is the sense in which we speak of *deprivation*: the point is not that AI produces poor text, but that it removes the productive struggle from which scholarly judgement grows.

Discussion of AI in the humanities has increasingly reflected this tension. Work on generative art asks whether authorship survives when a machine acts as a creative "muse" [3], and commentary on literature in the digital age asks how creativity is being reimagined under technological mediation [4]. These debates concern artists and writers, but their logic transfers directly to the academy. If authorship and creativity are contested when a machine contributes to a poem or an image, they are equally contested when a machine contributes the argument of a journal article.

1.2 Why this question is pressing now

Three lines of empirical work sharpen the concern. First, a large experimental study found that story ideas produced with generative AI assistance were rated as more creative at the individual level, particularly for writers who scored lower on baseline creativity, but that AI-assisted stories were more similar to one another [5]. Individual gains therefore coexisted with a collective loss of diversity. For a scholarly community whose progress depends on divergent thinking, this pattern deserves attention.

Second, research on cognitive offloading suggests that people readily delegate mental work to external tools and that delegation changes what they retain and practise. The "Google effect" studies showed that when people expect information to remain accessible, they remember where to find it rather than the information itself [8]. The broader offloading literature describes both the benefits of delegation and its cost to internal skill [9]. Applied to AI, a survey-based study of knowledge workers reported that greater confidence in AI was associated with less critical thinking effort, while greater confidence in one's own ability was associated with more [11]. A mixed-methods investigation of AI tool use likewise reported a negative association between frequent AI use and critical thinking, mediated by offloading [6]. An early electroencephalography preprint on essay writing reported weaker neural connectivity among participants using a language model than among those working unaided, although this finding is preliminary, based on a small sample, and not yet peer reviewed [7].

Third, educational evidence shows that assistance can help performance while harming learning. A field experiment with secondary-school mathematics students found that unrestricted access to a generative AI tutor improved performance on practice problems but reduced performance on later examinations taken without the tool, while a version designed

with pedagogical guardrails mitigated much of the harm [10]. The lesson is not that AI is harmful. It is that the *design of reliance* determines whether it builds or erodes capability.

1.3 The gap this paper addresses

The literature on AI in academia is dominated by two concerns: plagiarism and detection [13], and algorithmic bias [12]. Both are important, and both are treated by researchers within the supplied corpus [12,13]. Yet neither addresses the quieter loss this paper is concerned with. A scholar can use AI without plagiarising, without reproducing bias, and without any policy violation, and still gradually lose the habit of doing their own synthesis. This loss is invisible to detection software, absent from most integrity policies, and difficult to measure. It is also easy to overlook because the immediate output looks competent.

Moreover, the possibility of systemic degradation must be taken seriously. Language models are known to reproduce the statistical regularities of their training data rather than to understand meaning [14], and models trained recursively on synthetic output show progressive loss of the tails of the original distribution, a phenomenon termed model collapse [15]. If scholars increasingly draw their ideas from systems that favour the statistically typical, and if the scholarly record is increasingly populated with AI-mediated text, then the diversity on which creative recombination depends is at risk at both the individual and the systemic level.

1.4 Research questions

This review addresses four questions:

- **RQ1.** What mechanisms link reliance on AI to a reduction in creative synthesis among academics?
- **RQ2.** What does the available evidence, from both the supplied corpus and external anchor studies, show about each mechanism?
- **RQ3.** How do different disciplinary and ethical perspectives interpret the deprivation claim?
- **RQ4.** What safeguards are supported by the evidence?

1.5 Contribution and structure

The contribution is threefold. It offers a mechanism-based taxonomy of how AI reliance deprives academics of synthesis. It provides a transparent, reproducible screening of the supplied corpus alongside externally anchored evidence. And it provides a graded set of recommendations. Section 2 develops the conceptual framework. Section 3 reports the PRISMA-informed method. Sections 4 to 6 (Parts 2 and 3) present the related reviews and their results, the comparative analysis, and the multi-perspective appraisal. Sections 7 to 9 discuss implications, limitations, and conclusions.

2. Conceptual framework

2.1 What creative synthesis is

Creative synthesis is best understood as an *integrative act with four components*: retrieval of diverse material, tolerance of ambiguity, generation of unexpected associations, and evaluative

selection of those associations that cohere. Mednick's associative theory holds that creative individuals possess flatter associative hierarchies, so that remote elements become connectable, and that creativity is the formation of useful combinations from such elements [16]. Amabile's componential model adds that creative output depends not only on domain skill and creative-thinking skill but on task motivation, especially intrinsic motivation, which external rewards and constraints can undermine [17]. Boden's taxonomy [1] and Koestler's bisociation [2] complete the picture by stressing that novelty arises at the junction of separate conceptual spaces.

A particularly instructive scholarly example is Swanson's literature-based discovery. By connecting two bodies of literature that had never cited one another, he proposed that dietary fish oil might relieve Raynaud's syndrome, a hypothesis later supported clinically [18]. The lesson is that synthesis often requires a mind that has *personally* held two unrelated literatures in tension. A tool can retrieve both, but the recognition of a productive link depends on prior immersion, curiosity, and the felt sense that something is unresolved.

This immersion is what we mean by *formation*. The scholar who has read, struggled, misunderstood, and reread possesses a structured internal map of a field. It is that map, not access to text, that allows unexpected connections to surface. Delegating the reading and connecting to an external system may yield a summary, but not the map.

2.2 Five mechanisms of deprivation

Drawing on cognitive psychology, human factors research, and the wider literature, five mechanisms are proposed by which AI reliance may deprive academics of synthesis. These mechanisms organise the charting and analysis that follow.

Mechanism 1: Cognitive offloading. People shift mental tasks to external aids when the aids are available and convenient [9]. Offloading is often rational and can free capacity for higher-level work. The risk arises when what is offloaded is precisely the effortful integration that builds expertise. The memory literature shows that expected availability reduces retention of the content itself [8], and there is a plausible extension to conceptual work: if the connecting is always done externally, the internal connective structures are exercised less.

Mechanism 2: Loss of productive difficulty. Bjork and Bjork's concept of desirable difficulties holds that conditions which make learning feel harder in the short term, such as spacing, interleaving, and generating answers before seeing them, often improve long-term retention and transfer [19]. Generative AI is, by design, a difficulty-removal technology. It supplies the answer before the generation attempt. The mathematics field experiment discussed above is consistent with this account, since assisted practice performance rose while unassisted examination performance fell [10]. Synthesis, which depends on transfer across domains, is arguably the capacity most dependent on such difficulty.

Mechanism 3: Automation complacency and bias. Human factors research documents that operators who rely on automated aids monitor them less vigilantly and accept their outputs with excessive trust [20]. In scholarship, this translates into acceptance of a plausible AI-generated framing or citation without independent verification, and into a narrowing of the scholar's own evaluative habits. The survey evidence that higher confidence in AI accompanies lower critical thinking effort is consistent with this pattern [11].

Mechanism 4: Output homogenisation. Because language models tend toward high-probability continuations [14], and because recursive training on machine text erodes distributional tails [15], AI-mediated ideas tend toward the central and familiar. The finding that AI-assisted stories were individually better yet collectively more alike [5] is the clearest experimental illustration. For an academic field, homogenisation means fewer heterodox syntheses and more convergence around the already-consensual.

Mechanism 5: Epistemic narrowing. Models encode the perspectives that dominate their training corpora, and researchers in the supplied corpus have examined how AI can reproduce social stratification and bias [12]. Scholarship on decolonising sociology argues that indigenous and non-Western epistemologies have been marginalised in the production of theory [21]. Where AI systems are the default first reader of a literature, these marginal ways of knowing risk being further under-represented in the syntheses generated. Analyses of yogic and Vedantic epistemology likewise emphasise plural ways of knowing, including experiential and contemplative modes [22], which a text-prediction system is poorly placed to represent.

2.3 The attentional environment as an amplifier

AI reliance does not operate in a vacuum. Researchers in the supplied corpus have examined how digital media environments affect cognition, including the neuropsychology of social media use among adolescents [23], the cognitive overload and mental fatigue associated with doomscrolling [24], and digital minimalism as a form of resistance to attention capture [25]. These studies address populations and phenomena different from academic AI use, and they should not be read as evidence about it directly. Their relevance is contextual. They describe an attentional ecology in which sustained, effortful thought is already under pressure. In such an ecology, a tool that promises to reduce effort will be adopted quickly, and the countervailing habits of slow reading and reflective incubation will be harder to sustain.

2.4 An integrative model

The five mechanisms can be arranged as a causal chain. Convenience and time pressure encourage *substitutive* rather than *supplementary* use of AI. Substitutive use triggers offloading (M1), removes difficulty (M2), and fosters complacency (M3). Over time, these reduce the scholar's internal disciplinary map and evaluative sharpness. The scholar's outputs increasingly resemble the model's central tendencies (M4) and inherit its perspectival limits (M5). At the field level, the aggregate effect is a narrower and more homogeneous body of scholarship, which then feeds back into future training data. The model does not claim inevitability. It identifies the conditions, chiefly substitutive use without safeguards, under which deprivation becomes likely.

3. Methods

3.1 Design and reporting standard

A scoping review design was selected because the topic is emergent, the evidence is heterogeneous (experimental, survey, conceptual, and policy literature), and the aim is to map mechanisms and evidence rather than to estimate a pooled effect. Reporting followed the PRISMA 2020 flow logic [26] and the scoping-review extension PRISMA-ScR [27]. The use of PRISMA-ScR as a transparent method for mapping research landscapes is discussed in two

methodological treatments within the supplied corpus [28,29], which informed the charting approach adopted here.

3.2 Two evidence strands

Strand A: screening of the supplied corpus. The 117-entry reference list supplied for this study was treated as a defined corpus. It was screened at the level of bibliographic information (title, year, and venue), because abstracts and full texts were not available. This is a substantive limitation, addressed in Section 8, and it means that corpus records are cited for their stated subject scope rather than for specific quantitative findings.

Strand B: purposive anchor search. Because the corpus contains little experimental evidence on cognition and creativity, a purposive search identified foundational theoretical works and empirical studies on creativity, cognitive offloading, human-automation interaction, generative AI and learning, and language model limitations. Eighteen external sources were retained on the basis of relevance, peer-review status (with one clearly flagged preprint), and citation in the wider literature.

3.3 Eligibility criteria

Table 1. Eligibility criteria

Domain	Inclusion	Exclusion
Publication date	Published before 2026	Published in 2026 or later
Record completeness	Author, title, year, and venue identifiable	Missing year or incomplete bibliographic data
Topic	Addresses AI or algorithmic systems; creativity, authorship, or knowledge production; cognition, attention, or learning under digital conditions; research integrity or ethics; epistemology; or academic and educational systems	No plausible conceptual link to the research questions
Type	Empirical, theoretical, review, or policy analysis	Non-scholarly commentary without bibliographic attribution

3.4 Screening and selection (PRISMA-style flow)

Screening proceeded in four steps. First, records were checked for completeness and date eligibility. Second, duplicates were sought and none were found. Third, the remaining records were screened by title and venue against the topical criteria. Fourth, retained records were charted against the five mechanisms and against study type.

Table 2. PRISMA-informed flow of records

Stage	Strand A (supplied corpus)	Strand B (purposive anchors)
Records identified	117	18
Removed: publication year 2026	8	0
Removed: no publication date or incomplete record	2	0
Duplicates removed	0	0
Records screened	107	18
Excluded as topically unrelated	78	0
Records retained in synthesis	29	18
Total in synthesis	47	

The 78 excluded corpus records concerned subjects such as yoga physiology and philosophy, flood and climate studies, tort and contract law, accreditation processes, and public health. These are legitimate scholarly works but have no plausible bearing on AI reliance and synthesis. A further point of transparency is that the corpus is dominated by a single research group. This raises a risk of citation concentration, and it is the reason the anchor strand was added.

3.4 Data charting

Each retained record was charted for: year, venue type, study design, principal subject, relevance to each of the five mechanisms (direct, contextual, or none), and evidentiary weight. Evidentiary weight was graded in three tiers. **Tier 1** denotes peer-reviewed empirical evidence directly measuring an outcome relevant to synthesis or learning. **Tier 2** denotes peer-reviewed theoretical, conceptual, or review work. **Tier 3** denotes contextual or thematic works whose relevance is by analogy, together with any non-peer-reviewed preprint.

3.5 Integrity and originality safeguards

Because the paper concerns academic integrity, its own conduct must meet the standard it discusses. The manuscript was composed as original analytical prose. Source material was paraphrased and attributed, and direct quotation was avoided entirely. This practice reflects recognised guidance on plagiarism prevention and originality [30,31]. Ethical considerations relevant to social science research, including honesty in reporting, transparent method, and acknowledgement of limitations, were observed [32]. Authors should also submit the final manuscript to their institution's similarity-checking software, as no internal check can substitute for it.

3.6 Synthesis approach

Findings were synthesised narratively and by mechanism. No meta-analysis was attempted, because outcome measures, populations, and designs are too heterogeneous to pool, and because most corpus records are conceptual. Where evidence conflicted, the conflict is reported rather than resolved by selection.

4. Results: related research, its outputs, and what it establishes

4.1 Characteristics of the included evidence

Forty-seven records were retained (Table 2). Table 3 summarises them by strand, evidentiary tier, and design.

Table 3. Characteristics of included records (n = 47)

Category	Strand A: supplied corpus (n = 29)	Strand B: anchor sources (n = 18)
Tier 1: peer-reviewed empirical evidence	0 (not assessable at title level)	6
Tier 2: peer-reviewed theory, review, or conceptual work	0 (not assessable at title level)	9
Tier 3: contextual or thematic, or preprint	24 (provisional)	1
Methodological or integrity guidance	5	2
Publication years	2023 to 2025	1962 to 2025
Predominant design	Conceptual, thematic, or applied analysis	Experiments, surveys, reviews, theory

Two features of this profile should be stated plainly, because they define how much weight the conclusions can bear. First, no corpus record could be graded above Tier 3, because the grading requires access to the methods and findings, which were not available. This is a conservative rule and not a judgement of quality. Second, the strongest evidence bearing on the mechanisms comes from the anchor strand. Where the paper makes a causal or empirical claim, it therefore rests on the anchor studies, and where it draws on the corpus, it does so for context, framing, and conceptual convergence.

4.2 Anchor studies: designs and principal results

Table 4 sets out the empirical and computational anchor studies, their designs, what they found, and the mechanism to which each speaks. Results are reported at the level of direction and qualitative pattern.

Table 4. Principal empirical and computational anchor studies

Ref.	Design	Principal result	Mechanism	Tier
[5]	Online experiment: writers produced short stories with or without generative AI ideas; independent evaluators rated the stories	AI-assisted stories were rated as more creative and better written, with the largest gains among writers with lower baseline creativity; AI-assisted stories were more similar to one another	M4	1
[6]	Mixed-methods study (survey with follow-up interviews) of AI tool users	Higher AI use was associated with lower critical thinking scores, with cognitive offloading acting as a mediating pathway; younger participants showed greater dependence	M1, M3	1
[8]	Series of laboratory experiments on memory	When people expected to access information later, they recalled the information less well and recalled where to find it better	M1	1
[10]	Large field experiment in secondary-school mathematics with a guarded and an unguarded AI tutor	Unrestricted AI access raised practice performance but lowered later unassisted examination performance; the guarded tutor largely mitigated the harm	M2	1
[11]	Survey of knowledge workers reporting how they use generative AI	Confidence in AI was associated with less critical thinking effort; confidence in one's own ability was associated with more	M3	1
[15]	Computational experiments on recursive training of generative models	Models trained repeatedly on machine-generated output progressively lost the rare, low-probability content of the original data distribution	M4	1
[7]	Small-sample EEG study of essay writing under three conditions (language model, search engine, no tools)	Participants using the language model showed weaker neural connectivity and weaker recall of their own essays; findings are preliminary and not peer reviewed	M1, M2	3

Three observations follow from the table. The first is the *heterogeneity of measures*. Creativity ratings, self-reported effort, recall, exam performance, and distributional statistics are not

commensurable, which is why no pooled estimate was attempted. The second is *convergence of direction* despite that heterogeneity. Across memory, learning, effort, and creative diversity, the same shape recurs: assistance improves an immediate product or performance while reducing an internal or collective quantity. The third is that *none of the studies measures the creative synthesis of practising academics directly*. Every study either uses a proxy task (story writing, mathematics problems, essay writing) or a self-report. This is the central evidence gap of the field and is developed in Section 5.

The theoretical anchors provide the interpretive scaffolding. Boden's taxonomy [1], Koestler's bisociation [2], Mednick's associative theory [16], and Amabile's componential model [17] define what synthesis is. Bjork and Bjork [19] explain why difficulty matters to durable learning. Risko and Gilbert's review [9] and Parasuraman and Manzey's review [20] supply the general accounts of offloading and complacency, and Bender and colleagues [14] supply the critique of what language models do and do not do. Swanson's literature-based discovery [18] provides the paradigmatic case of scholarly synthesis.

4.3 The supplied corpus: thematic results

The 29 corpus records fall into seven themes. Table 5 presents the themes, the records, and the role each theme plays in the argument. Records are described by stated subject scope.

Table 5. Thematic distribution of retained corpus records (n = 29)

Theme	Records	Stated scope	Role in the argument
T1. Creativity, authorship, digital culture	[3], [4], [33], [34], [35]	Generative art and authorship; creativity in literature; changing value of art in a post-digital economy; memes as cultural capital; self-branding and identity	Frames creativity and authorship as socially contested under digital mediation
T2. AI ethics, bias, labour	[12], [36], [37], [38]	Algorithmic bias and stratification; digital surveillance and sovereignty; algorithmic invisibility among gig workers; work ethic and alienation	Supports M5 (epistemic narrowing) and the political economy of automation
T3. Integrity and method	[13], [30], [31], [32], [28], [29]	AI and plagiarism in publishing; detection tools; originality practice; research ethics; PRISMA-ScR	Provides the integrity and method context
T4. AI in professional practice	[39], [40], [41], [42], [43]	AI in diagnosis; AI-driven patient engagement; machine learning in facility management; AI-based prediction of antimicrobial resistance; digital	Illustrates AI adoption in expert professional settings as a comparator

		transformation in hospital management	
T5. Education and research systems	[44], [45], [46]	Technology programmes for education; drawbacks of national education policy; accreditation and research promotion	Frames the institutional incentives surrounding academic work
T6. Cognition and attention	[23], [24], [25], [47]	Social media and neuropsychology; cognitive overload; digital minimalism; brain plasticity and practice	Supplies the attentional ecology and the plasticity premise
T7. Epistemology	[21], [22]	Indigenous epistemologies in sociology; ways of knowing in yogic philosophy	Supports M5 through the argument for epistemic plurality

T1: Creativity, authorship, and digital culture. These records show that the meaning of creativity and authorship is already unsettled in the arts. The treatment of AI as a creative muse [3] and the reflection on literary creativity in the digital age [4] address authorship directly. Work on the changing value of art after non-fungible tokens [33], on memes as cultural capital [34], and on self-branding in social media [35] add the point that digital environments reward certain kinds of cultural production, typically the quickly legible, shareable, and reproducible. This matters for the present argument because the conditions in which scholars now work are themselves shaped by metrics of visibility and speed, and AI tools fit those conditions well.

T2: AI ethics, bias, and labour. The analysis of algorithmic bias and social stratification [12] establishes that AI systems are not neutral mirrors of knowledge. The studies of digital surveillance and sovereignty [36], of algorithmic invisibility among gig workers [37], and of work ethic and alienation in modern employment [38] describe how automated systems reorder who is seen, who is rewarded, and how workers relate to their own labour. Read together, they suggest an important analogy: academic work is increasingly platform-mediated and metric-governed, and the alienation from one's own work that these studies describe in other sectors has a plausible academic counterpart when synthesis is delegated.

T3: Integrity and method. The corpus is strongest here. It addresses the use of AI to combat plagiarism in publishing [13], the tools and techniques of plagiarism detection [30], practices for producing original content [31], and the ethics of social science research [32]. Two treatments of PRISMA-ScR [28,29] supply method. This body of work is directly relevant but conceptually adjacent: it is concerned with *what is written and whether it is original*, not with *what the writer becomes by writing*. That distinction is the core of the gap.

T4: AI in professional practice. The corpus includes applications of AI to diagnostic medicine [39], patient engagement [40], smart facility management [41], prediction of antimicrobial resistance [42], and the digital transformation of hospital management [43]. These are not studies of academic synthesis, and they are not cited as evidence about it. They are useful as comparators, because in each of these domains a professional must decide how far to defer to

an algorithmic output. That question, how far expertise should defer, is the same question that faces a scholar deciding whether to accept an AI-generated framing.

T5: Education and research systems. Records on technology programmes for education [44], on the drawbacks of a national education policy [45], and on the role of accreditation in promoting research and innovation [46] describe the institutional environment. Systems that reward measurable output, whether accreditation indicators or publication counts, create pressure toward the fastest route to output. This is a structural driver of substitutive AI use and is developed in Section 6.

T6: Cognition and attention. The neuropsychological analysis of social media use [23], the study of cognitive overload and mental fatigue [24], and the treatment of digital minimalism [25] describe an attentional ecology under strain. A study of how practice reshapes the brain [47] contributes the general premise that repeated activity structures neural organisation, although it concerns a different practice and is cited here for the principle only. If practice builds capacity, then the removal of practice may plausibly fail to build it. This is an inference, not a demonstrated result.

T7: Epistemology. The argument for decolonising sociological theory through indigenous epistemologies [21] and the analysis of ways of knowing in yogic philosophy [22] support a pluralist view of knowledge. They ground the concern that a system trained on dominant corpora will under-represent alternative ways of knowing, which is the substance of Mechanism 5.

4.4 Related reviews and analyses: outputs and limits

The task of situating this paper requires an honest account of what prior reviews exist. No review or meta-analysis specifically on AI reliance and academic creative synthesis was identified in either strand. The nearest neighbours are reviews and analyses in adjacent domains. Table 6 sets out those neighbours, what each contributes, and what each leaves unanswered.

Table 6. Nearest related reviews and analyses: outputs and unanswered questions

Ref.	Type	Output	What it leaves unanswered
[13]	Analysis (scholarly publishing)	Examines AI as a tool for combating plagiarism	Does not address whether AI use, even when originality is preserved, changes the scholar's own capacities
[3]	Conceptual analysis (arts)	Reexamines artistic authorship when AI acts as a creative collaborator	Concerns artistic rather than scholarly synthesis; not an empirical test
[4]	Conceptual analysis (literature)	Reflects on creativity in the digital age	Does not isolate AI reliance as a variable

[43]	Systematic review (hospital management)	Reviews the impact of digital transformation on a professional field	A different domain; does not address creativity or cognition
[28, 29]	Methodological treatments	Provide procedures for transparent scoping reviews	Method only; no substantive finding on AI
[9]	Review (cognitive science)	Synthesises how and why people offload cognition	Predates generative AI; does not address creative synthesis
[20]	Review (human factors)	Synthesises evidence on complacency and bias in automation use	Concerns monitoring tasks, not idea generation
[6]	Mixed-methods study	Links AI use to critical thinking through offloading	Measures critical thinking, not synthesis; correlational

The pattern is clear. The literature contains rich work on *plagiarism*, *bias*, *offloading*, and *learning*, but the specific question of whether reliance on AI erodes the integrative capacity of scholars has been approached only indirectly. This finding justifies the present review as a mapping exercise and justifies caution in the strength of its conclusions.

5. Comparative analysis

5.1 Mechanism-by-evidence comparison

Table 7 compares the five mechanisms on the evidence available, the directness of that evidence to academic synthesis, and the principal counter-evidence.

Table 7. Comparative appraisal of the five mechanisms

Mechanism	Core supporting evidence	Directness to academic synthesis	Principal counter-evidence or boundary	Confidence
M1. Cognitive offloading	[6,8,9]; preliminary [7]	Indirect (memory and critical-thinking proxies)	Offloading can release capacity for higher-order work [9]	Moderate
M2. Loss of productive difficulty	[10,19]; preliminary [7]	Indirect (learning outcomes in students)	Guarded tool design largely removed the harm [10]	Moderate
M3. Automation complacency	[11,20]	Indirect (knowledge workers; automation tasks)	Self-confidence appears protective [11]	Moderate

M4. Output homogenisation	[5,14,15]	Partial (creative writing task; model-level dynamics)	Individual creativity ratings improved [5]	Moderate to strong for the pattern; limited for scholarship
M5. Epistemic narrowing	[12,14,21,22]; context [36,37]	Conceptual	Little direct empirical test on scholarly outputs	Limited

The comparison exposes an asymmetry. The mechanisms with the strongest experimental backing (M4, M2) are supported by studies of tasks that only resemble scholarly synthesis, and the mechanism most specific to scholarship (M5) has the thinnest empirical base. A rigorous reader should therefore treat the mechanisms as *well-motivated hypotheses with convergent partial support*, not as established facts about academics.

5.2 Assisted performance versus unassisted capability

A recurring analytical distinction concerns what is measured while the tool is present versus what remains when it is removed. Table 8 recasts the anchor findings on this axis.

Table 8. Assisted versus unassisted or collective outcomes

Ref.	Outcome with assistance	Outcome without assistance, over time, or at collective level	Direction of divergence
[10]	Practice performance rose with unrestricted AI	Later unassisted examination performance fell	Assisted gain, unassisted loss
[5]	Individual stories rated more creative	Collective diversity of ideas fell	Individual gain, collective loss
[8]	Information readily accessible	Recall of content itself was weaker	Access gain, retention loss
[11]	Task completed with generative AI	Confidence in AI accompanied less critical effort	Efficiency gain, effort loss
[15]	Model output produced at scale	Rare content disappeared over generations of retraining	Volume gain, diversity loss

Every row displays the same structure. This does not prove a common cause, since the studies differ in design and quality. It does, however, support a general analytical principle: *measures taken during assisted performance systematically overstate capability, and measures taken at the individual level systematically overstate collective health*. For universities that assess the productivity of researchers by output volume, this is a serious measurement problem. An academic whose output rises after adopting AI may be more productive, or may be less formed, and the indicator cannot distinguish them.

5.3 Who is most at risk: expertise as a moderator

The anchor evidence contains a subtle but important signal. In the creative writing experiment, the benefit of AI assistance was greatest for writers with lower baseline creativity [5]. In the mathematics experiment, students without guardrails were the population harmed [10]. If we apply these patterns to scholarship with appropriate tentativeness, two implications follow.

First, *novices may gain the most in the short term and lose the most in the long term*. An early-career researcher who has not yet built a disciplinary map may find AI assistance an enormous relief, and may therefore be least likely to develop the map. An experienced scholar, who already possesses the internal structure, may use AI to accelerate work without losing the capacity that makes acceleration meaningful. This is the logic of Amabile's componential model [17]: domain expertise and creative-thinking skill are inputs to creativity, and a tool that substitutes for their acquisition undermines the inputs.

Second, the effect of AI is likely *conditional on the task*. Use for mechanical operations (formatting references, checking grammar, translating a passage) is unlikely to touch synthesis. Use for the generation of research questions, the framing of theoretical positions, and the interpretation of contradictory findings is where substitution is most consequential. The classification of tasks by their proximity to synthesis is therefore an essential design principle and is developed in Section 8.

5.4 Institutional and cultural comparators

The corpus offers a set of comparators that place the individual scholar in an institutional setting. Research on the promotion of research and innovation through accreditation [46], on the drawbacks of national education policy [45], and on technology programmes in education [44] each concern systems that shape what academics are rewarded for. Systems that reward measurable output over depth generate demand for fast production. Studies of work ethic and alienation [38] and of platform-mediated labour [37] describe how workers in other sectors respond to comparable pressures: by minimising effort where effort is unrecognised, or by adapting to metrics that render deep work invisible. The academic analogue is the researcher who uses AI to satisfy an output metric without the intellectual investment the metric was meant to signal.

The comparison also cautions against a purely individualist reading. If deprivation of synthesis occurs, it is not only a failing of individual discipline. It is partly an *institutional design problem*, in which incentives, workloads, and evaluation practices make substitutive AI use the rational choice.

5.5 Comparison of perspectives on the same evidence

The same evidence supports different readings depending on the theoretical lens. Table 9 summarises how five perspectives interpret the deprivation claim. These perspectives are developed further in Part 3.

Table 9. Perspectives on the deprivation claim

Perspective	Reading of the evidence	Central emphasis
Cognitive psychology	Offloading and lost difficulty reduce internal capacity [8,9,19]	Process and practice build capability
Creativity studies	Homogenisation narrows the pool of novel combinations [5,16,17]	Diversity and remote association
Human factors	Complacency erodes evaluative vigilance [11,20]	Design of human-machine interaction
Critical and postcolonial theory	Models reproduce dominant epistemologies [12,14,21]	Whose knowledge is amplified
Technological optimism	Guarded design can preserve or enhance capability [5,10]	Augmentation, not substitution

No single perspective is sufficient. The optimistic reading is supported by the finding that individual creativity ratings rose [5] and that guardrails reduced harm [10]. The pessimistic reading is supported by the collective and unassisted deficits recorded in the same studies. The most defensible synthesis is a *conditional* one: AI's effect on synthesis is determined by the pattern of use, not by the technology alone.

5.6 Citable analytical propositions

The comparative analysis yields six propositions that are offered as hypotheses for empirical test. They are presented as this paper's analytical contribution and are intended to be citable as such, not as established findings.

- **P1 (Process-product decoupling).** In AI-assisted scholarship, the quality of the product is an unreliable indicator of the development of the producer.
- **P2 (Proximity gradient).** The risk of deprivation increases with the proximity of the delegated task to the integrative core of synthesis (question framing, connection-making, interpretation) and is negligible for clerical tasks.
- **P3 (Novice vulnerability).** Researchers who have not yet built an internal disciplinary map are at greater long-term risk from substitutive AI use than established researchers, although they may perceive the greatest short-term benefit.
- **P4 (Collective-individual divergence).** Field-level diversity of ideas may decline even as individual output quality rises.
- **P5 (Institutional amplification).** Output-metric evaluation systems amplify substitutive AI use by rewarding volume over formation.
- **P6 (Design dependence).** The effect of AI on synthesis depends on interaction design; tools that require the scholar to generate before receiving assistance preserve more capability than tools that supply answers first.

5.7 Boundary conditions and counter-arguments

A critical analysis must state where the deprivation claim does not hold. First, offloading is not inherently harmful. The offloading literature stresses that delegation can free cognitive resources for more valuable activity [9], and there are real cases in which AI's retrieval and summarisation capacity allows a scholar to survey far more literature than unaided reading would permit. Swanson-style discovery was originally a manual process of connecting disjoint literatures [18], and the difficulty of that process is arguably one reason such discoveries were rare. Tools that widen the scholar's field of view could increase the number of potentially connectable literatures.

Second, the harms documented in the anchor studies were not uniform. The guarded tutoring condition [10] shows that deliberate design can preserve learning, and the finding that self-confidence in one's own ability predicted more critical effort [11] suggests that a scholar's stance towards AI matters. Third, some evidence is preliminary. The most striking neural findings come from a small preprint [7], and the correlational studies [6,11] cannot establish causation. It is possible that people who already think less critically are more likely to rely on AI, rather than the reverse.

These counter-arguments do not dissolve the concern. They refine it: the risk is real but conditional, and the appropriate response is targeted design and empirical monitoring, not prohibition and not complacent adoption.

6. Multi-perspective analysis

A claim as broad as "AI reliance deprives academics of creative synthesis" cannot be judged from one disciplinary standpoint. This section examines it from seven perspectives. The aim is to test the claim against the strongest reading each perspective would give it, including readings that weaken it.

6.1 The cognitive-psychological perspective

Cognitive psychology treats synthesis as the output of internal structure built through practice. On this view, what matters is the *history of effortful processing* behind an insight, not the insight alone. Expected availability of information alters what is stored [8]. Delegating cognition changes the balance between internal skill and external support [9]. Conditions that make learning feel harder often improve durable learning and transfer [19].

This perspective's strongest contribution is the distinction between *performance* and *learning*. A scholar working with an AI assistant performs well, and this performance is visible and measurable. Whether the scholar is learning is a separate question, and the mathematics field experiment showed that the two can diverge sharply [10]. Synthesis depends on transfer, the ability to bring knowledge from one domain to bear on another, so it is arguably the capacity most exposed to this divergence.

The perspective also supplies the main caution. The neural evidence in the anchor set is preliminary [7], and the largest survey-based associations are correlational [6,11]. The perspective supports the hypothesis of deprivation, but it does not yet confirm it for academics specifically.

6.2 The creativity-studies and arts perspective

From the standpoint of creativity research, deprivation is chiefly a problem of *diversity*. Mednick's associative account holds that novelty arises from remote associations [16]. Boden's combinational creativity and Koestler's bisociation both require that distant frames be brought together [1,2]. A system that tends toward high-probability continuations [14] will tend to propose the more proximate associations. The finding that AI-assisted stories were individually stronger but collectively more alike [5] is a direct empirical statement of this concern.

The arts-and-humanities literature in the corpus extends the point from *ideas* to *authorship*. The reconsideration of authorship when AI acts as a creative muse [3] and the reflection on creativity in digital literature [4] both suggest that creativity is not only a property of outputs but a relationship between a maker and a work. The analysis of how the value of art changes in a post-digital economy [33] and of memes as cultural capital [34] add that creative value is socially negotiated and shaped by the media through which work circulates. Applied to scholarship, this implies that a synthesis is valued partly because a *particular mind* produced it, and that the erosion of this link changes what scholarship is, not only how much of it there is.

A fair counterpoint from this perspective is that creative traditions have always used tools, collaborators, and chance procedures, and that authorship has never been purely solitary. Individual creativity gains from AI assistance were real in the experimental evidence [5], and gains were largest for those with less baseline creativity. The perspective therefore allows a democratising reading, in which AI raises the floor of creative competence, alongside the homogenising reading, in which it lowers the ceiling of collective novelty. Both readings can be true at once, at different levels.

6.3 The pedagogical and early-career perspective

Pedagogy adds a developmental dimension. Doctoral training is a formation process in which students become scholars partly by struggling with literature, methods, and argument. Amabile's model treats domain skill and creative-thinking skill as prerequisites of creative work [17]. If early-career researchers use AI for the very tasks through which those skills are built, they may produce competent text without acquiring the underlying competence.

This perspective helps explain why the deprivation may be *invisible to the person experiencing it*. The complacency literature shows that trust in automation can reduce vigilance [20], and the knowledge-worker survey shows that confidence in AI accompanies less critical effort [11]. A doctoral student who has never synthesised a literature unaided has no baseline against which to notice what has been lost. Institutional documents also show the educational system under strain. The analysis of drawbacks in national education policy [45] and the review of technology programmes in education [44] indicate that educational reform and educational technology are both proceeding rapidly and that their consequences for depth of learning deserve scrutiny.

The pedagogical response most consistent with the evidence is not prohibition but *sequencing*. The guarded tutoring design [10] worked because it changed the order of interaction, requiring the learner to attempt before being helped. Doctoral supervision could adopt the same logic: unaided synthesis first, AI critique second.

6.4 The research-integrity and ethics perspective

The corpus is at its strongest on integrity. It addresses AI as a tool for combating plagiarism in scholarly publishing [13], the detection techniques and tools available to editors [30], and practices for producing original content [31]. Ethical guidance for social science research stresses honesty, transparency, and responsibility in method [32].

This perspective expands the meaning of integrity. Conventional integrity frameworks ask whether the text is original and properly attributed. They rarely ask whether the *thinking* behind the text belongs to the named author. A manuscript can pass a similarity check, cite every source correctly, and still represent an argument the author did not really work out. This is a form of intellectual ghost-writing by a machine, and existing detection tools are not designed to identify it. The integrity question shifts from *what was copied* to *what was understood*.

The perspective also warns against an integrity approach that relies on detection alone. Detection tools examine text properties, whereas the deprivation problem concerns cognitive process. A policy that focuses on detection will therefore miss the harm this paper describes and may generate false confidence. A better integrity architecture would combine disclosure of AI use, process documentation, and oral or dialogic assessment of understanding.

6.5 The institutional and political-economy perspective

Political economy asks why scholars would choose substitutive AI use, and it locates much of the answer in incentives. Accreditation regimes and research-promotion frameworks reward measurable output [46]. Labour studies show how workers under metric-driven, platform-mediated systems adapt in ways that serve the metric rather than the work's purpose. This includes the alienation described in studies of quiet quitting [38] and the invisibility that algorithmic systems impose on workers whose effort is not captured by the platform [37]. Analyses of digital surveillance and sovereignty [36] add that technological systems are embedded in structures of power that determine who controls the tools and the data.

The academic implication is that AI adoption is not a neutral personal choice. When universities rank staff on publication counts, when funders reward rapid proposals, and when journals face rising submission volumes, AI becomes a rational response to overload. On this analysis, the blame for any deprivation of synthesis lies substantially with system design. It also follows that individual exhortations to "use AI responsibly" are likely to be weak unless the surrounding incentives change.

6.6 The epistemic-justice and decolonial perspective

This perspective asks whose knowledge is amplified. Analysis of algorithmic bias shows that AI can reproduce and deepen social stratification [12]. The critique of large language models identifies risks from training on large, uncured corpora that over-represent dominant voices [14]. The decolonial scholarship in the corpus argues that indigenous epistemologies have been marginalised in the making of sociological theory [21], and the analysis of yogic epistemology describes a plural account of knowing that includes contemplative and experiential modes [22].

The relevance to synthesis is that *synthesis is not epistemically neutral*. Choosing which literatures to connect is a decision about whose knowledge counts. If scholars ask AI systems

to summarise "the literature" or to propose "the gap", the system's implicit weighting of sources becomes the scholar's default. Scholars working at the margins of dominant paradigms, or in languages and traditions that are under-represented in training data, may find that AI-mediated synthesis pulls them towards the mainstream. The concern is aggravated by model collapse dynamics, in which recursive training erodes the tails of the distribution [15], and the tails are precisely where marginal knowledge lives.

A caution is warranted. The corpus works in this area are conceptual and cited for their argument, not for measured effects on AI-assisted scholarship. The claim that AI narrows scholarly perspectives is plausible and theoretically well motivated, but it remains largely untested.

6.7 The professional-practice comparator

Professional fields that have adopted AI earlier offer a useful comparison. The corpus includes work on AI in diagnostic medicine [39], patient engagement in AI-driven healthcare [40], machine learning for facility management [41], AI prediction of antimicrobial resistance [42], and the effects of digital transformation on hospital management [43]. These studies address technical performance and organisational effects, not creativity, and they are not cited as evidence about synthesis.

Their value is as a *structural analogy*. In diagnostic settings, the core professional question is how far the expert should defer to an algorithm, and the human-factors literature on complacency [20] treats that question as central to safe design. Medicine has developed norms of verification, second reading, and accountable sign-off precisely because unexamined deference is dangerous. Academia has developed no equivalent norm for AI-mediated synthesis. A scholar who accepts an AI-generated framework without independent construction resembles a clinician who accepts an algorithmic diagnosis without examining the patient. The analogy suggests that scholarly practice needs its own verification norms.

6.8 The augmentation (optimist) perspective

The strongest opposing view holds that AI is an amplifier and that deprivation is a design failure, not a necessity. Its evidence includes the improvements in individual creativity ratings [5], the mitigation achieved by guarded tutor design [10], and the observation that offloading can release capacity for higher-order tasks [9]. Its historical argument is that every new cognitive tool, from writing to calculators to search engines, provoked fears of intellectual decline, and that scholarship adapted and often improved. The memory literature itself shows that people adapt to external stores by reorganising what they retain [8], not simply by losing capacity.

This perspective deserves its full weight, and this paper's conclusions incorporate it through the conditional framing of the propositions. Its limitation is that historical analogies are imperfect. Earlier tools offloaded storage, calculation, or retrieval. Generative models offload *composition and interpretation*, the very activities in which scholarly identity is formed. The claim that this difference is decisive is a hypothesis, and the claim that it is irrelevant is equally one.

6.9 Summary of perspectives

Table 10. Synthesis of perspectives

Perspective	Verdict on deprivation	Key condition	Key reference points
Cognitive psychology	Plausible; partly supported	Whether effortful processing is replaced	[7-10,19]
Creativity and arts	Supported at collective level; contested at individual level	Level of analysis	[1-5,16,33,34]
Pedagogy and early career	Highest risk for novices	Sequencing of unaided and assisted work	[10,17,44,45]
Integrity and ethics	Existing frameworks miss the harm	Focus on process, not only text	[13,30-32]
Institutional and political economy	Substantially structural	Incentive design	[36-38,46]
Epistemic justice	Plausible; largely untested	Source diversity of the tool	[12,14,15,21,22]
Augmentation	Deprivation is avoidable	Interaction design	[5,9,10]

Across the seven perspectives, agreement is strongest on one point: *outcomes depend on how AI is used, by whom, and under what incentives*. The perspectives disagree mainly about how likely harmful patterns of use are to prevail and about whether the harms are reversible.

7. Discussion

7.1 Principal findings

This review set out to determine how AI reliance may deprive academics of creative synthesis and how strong the evidence is. Five findings stand out.

First, **the mechanisms are well motivated and convergently supported, but not directly demonstrated for academics**. Cognitive offloading, loss of productive difficulty, complacency, homogenisation, and epistemic narrowing each have support from experimental, survey, or theoretical work [5-11,14,15,19,20]. Yet no study measured the integrative work of practising scholars directly. The claim that AI deprives academics of synthesis is therefore best described as a strongly supported hypothesis, not an established fact.

Second, **the same divergence recurs across designs**. In every anchor study with both an assisted and an unassisted or collective measure, the assisted or individual measure improved while the unassisted or collective measure deteriorated or did not follow (Table 8). This recurring structure is the most informative pattern in the evidence and the basis of Proposition P1.

Third, **the deprivation is conditional and design-dependent**. The guarded-tutor result [10] and the protective role of self-confidence [11] show that outcomes vary with interaction design and user stance. This is the finding with the greatest practical importance, because it implies that the problem is addressable.

Fourth, **institutions amplify the risk**. Output-metric systems [46] and platform-style labour dynamics [37,38] make substitutive use economically rational, so responsibility cannot rest on individuals alone.

Fifth, **the integrity discourse has a blind spot**. The literature on plagiarism and detection is substantial [13,30,31], but it addresses the originality of the *text*, not the formation of the *thinker*. A scholar can be fully compliant and still be deprived.

7.2 Interpreting the process-product problem

The organising insight of this paper can be stated as a distinction between *production* and *formation*. AI systems are extraordinarily effective at production. They are, by design, indifferent to formation. Scholarship, however, has two outputs: the work and the scholar. Universities and journals evaluate the work, while society ultimately depends on the scholar, because it is the trained scholar who can evaluate claims, notice errors, and recognise which connection is worth pursuing. An evaluation system that measures only the work is blind to the second output.

Swanson's example [18] makes the point concrete. The value of his discovery lay not in the retrieval of two literatures, which any modern tool can do, but in the judgement that a particular connection was meaningful. That judgement rested on his prior immersion. If a generation of scholars is trained to receive connections rather than to make them, the supply of such judgement may thin, even as the supply of plausible connections increases. This is the paradox at the centre of the deprivation claim: *abundant synthesis-like output may coexist with scarce synthesising capacity*.

7.3 Relation to the existing debate

The current debate on AI in scholarship is dominated by questions of permissibility (is AI use allowed?), detection (can it be identified?), and disclosure (must it be declared?). These are necessary questions but treat AI as a threat to the *integrity of the record*. This paper reframes it as a possible threat to the *capacity of the community*. The difference has policy consequences. Integrity threats are addressed by rules and enforcement. Capacity threats are addressed by training design, evaluation reform, and cultural norms that protect the conditions for deep work. The attentional-ecology literature in the corpus [23-25] suggests that such protective norms are already under pressure from other digital forces, which is a reason for urgency.

7.4 Theoretical contributions

The paper contributes a mechanism taxonomy (M1 to M5) that links the cognitive, creative, human-factors, and epistemic literatures, and an integrative causal model in which substitutive use is the trigger condition (Section 2.4). It contributes six propositions (P1 to P6) as testable hypotheses. It also contributes the process-product distinction as an analytical lens that can be applied to other domains of AI-assisted professional work, including the clinical and managerial fields represented in the corpus [39-43].

8. Recommendations and limitations

8.1 Recommendations

The recommendations follow from Propositions P2, P5, and P6 and are graded by the strength of the supporting evidence. They aim to preserve the benefits of AI while protecting the conditions for synthesis.

Table 11. Evidence-graded recommendations

Level	Recommendation	Rationale	Evidence grade
Individual scholar	Generate before you consult. Draft your own research question, argument outline, or interpretation before asking AI to critique or extend it	Preserves generation and difficulty [10,19]	Moderate (proxy tasks)
Individual scholar	Classify tasks by proximity to synthesis. Delegate clerical tasks freely; treat framing, connection-making, and interpretation as protected	Follows P2; offloading is benign for low-level tasks [9]	Moderate (conceptual)
Individual scholar	Verify actively. Independently check AI-supplied claims, citations, and framings	Counters complacency [11,20]	Moderate
Individual scholar	Protect deep-work time. Schedule sustained unassisted reading and reflection	Counters attentional pressure [23-25]	Limited (contextual)
Supervisors	Sequence doctoral training. Require unaided literature synthesis at defined stages before AI-assisted refinement	Protects formation of the disciplinary map [17]	Moderate
Supervisors	Assess understanding, not only text. Use oral defence, process logs, and iterative drafts	Addresses the integrity blind spot [13,30,31]	Limited to moderate

Institutions	Reform indicators. Reduce weight on output volume; reward depth, replication, and integrative work	Addresses institutional amplification, P5 [46]	Limited (conceptual)
Institutions	Adopt process-aware integrity policies. Combine disclosure, documentation, and dialogic assessment rather than detection alone	Detection cannot capture process [30,32]	Limited to moderate
Tool designers	Build friction by design. Prompt users to attempt first, show reasoning, and offer critique in place of answers	Guardrails reduced harm [10]	Moderate
Tool designers and researchers	Preserve diversity. Audit tools for source diversity and monitor homogenisation	Addresses M4 and M5 [5,12,15]	Limited
Funders and journals	Fund direct studies of AI use and synthesis among practising academics	Fills the central evidence gap	Not applicable

8.2 A research agenda

Five studies would substantially improve the evidence base. The first is a **longitudinal cohort study** of doctoral students with differing patterns of AI use, measuring unaided synthesis performance over time. The second is a **randomised design** comparing "generate-first" with "answer-first" AI workflows on the originality and diversity of research proposals. The third is a **bibliometric analysis** of idea diversity, such as conceptual and citation diversity, in fields before and after widespread AI adoption. The fourth is a **cross-cultural comparison** of AI use and synthesis in linguistic and epistemic traditions that are differently represented in training data, which would test M5 directly. The fifth is a **validated instrument** for measuring integrative capacity, since none currently exists for scholarly synthesis. Each of these would convert a proposition in this paper into an empirical result.

8.3 Limitations

This review has limitations that readers should weigh carefully.

- **Bibliographic-level screening of the corpus.** The supplied corpus was screened by title, year, and venue, and abstracts and full texts were not available. Corpus records are therefore cited for their stated scope and not for specific findings, and they could not be graded above Tier 3. Some records may be more or less relevant than their titles suggest.
- **Concentration of the corpus.** The corpus is dominated by a single research group and a small set of venues. This limits independence of evidence and could introduce citation concentration. The purposive anchor strand was added to offset this, but it is not a systematic database search.

- **No systematic database search.** Strand B was purposive, not exhaustive. Relevant studies may have been missed, and the review does not claim completeness. The PRISMA-informed structure serves transparency, not full systematic-review rigour.
- **Proxy evidence.** No included study measured the synthesis of practising academics directly. Story-writing, mathematics learning, essay-writing, and survey outcomes are proxies, and generalisation to scholarship is an inference.
- **Correlational and preliminary evidence.** Some key findings are correlational [6,11] or preliminary [7], and reverse causation cannot be excluded.
- **Temporal limit.** The field is moving quickly. Only sources published before 2026 were included, and later evidence may alter the picture.
- **Absence of new primary data.** No new empirical data were collected. The comparative tables are syntheses of existing sources, and the propositions are hypotheses and not results.
- **Selection judgement.** Screening decisions involve judgement, which was applied transparently but was not independently replicated by a second reviewer.

9. Conclusion

The claim that AI reliance deprives academics of creative synthesis is neither a moral panic nor a settled fact. It is a well-motivated hypothesis, supported by converging evidence from cognitive psychology, creativity research, human-factors science, and the critical study of language models, yet not directly tested on the population it concerns.

The central argument of this paper is that scholarship yields two products, the work and the scholar, and that AI systems are optimised for the first while being indifferent to the second. When AI is used as a substitute for the effortful integration through which scholars are formed, output can rise while capacity quietly declines. This decoupling is invisible to the metrics that universities and publishers currently use, and it is invisible to detection tools that examine text, not process. At the collective level, tools that favour the statistically typical, and that are increasingly trained on their own outputs, may narrow the diversity of ideas on which recombination depends.

The evidence also shows that this outcome is not inevitable. Interaction design that requires generation before assistance, awareness of one's own confidence and vigilance, protection of tasks close to the integrative core of scholarship, and institutional reform of evaluation incentives can all preserve the conditions for synthesis. The choice is not between using AI and refusing it. It is between using it in ways that extend the scholar's reach and using it in ways that replace the scholar's mind.

The task for the academic community is therefore threefold: to study the question directly, to redesign its incentives so that depth is recognised, and to cultivate norms in which the struggle of synthesis is treated as the substance of scholarship and not as an inefficiency to be removed.

References

1. Boden MA. The creative mind: myths and mechanisms. 2nd ed. London: Routledge; 2004.
2. Koestler A. The act of creation. London: Hutchinson; 1964.
3. Ahluwalia PS, Khandey PD. AI as muse: redefining artistic authorship in the age of generative art. *Siddhanta's Int J Adv Res Arts Humanit.* 2025:15-33.

4. Kumari B. Reimagining creativity: literature in the digital age. *Siddhanta's Int J Adv Res Arts Humanit.* 2025:25-32.
5. Doshi AR, Hauser OP. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Sci Adv.* 2024;10(28):eadn5290.
6. Gerlich M. AI tools in society: impacts on cognitive offloading and the future of critical thinking. *Societies.* 2025;15(1):6.
7. Kosmyna N, Hauptmann E, Yuan YT, Situ J, Liao XH, Beresnitzky AV, et al. Your brain on ChatGPT: accumulation of cognitive debt when using an AI assistant for essay writing task. *arXiv [Preprint].* 2025:2506.08872.
8. Sparrow B, Liu J, Wegner DM. Google effects on memory: cognitive consequences of having information at our fingertips. *Science.* 2011;333(6043):776-8.
9. Risko EF, Gilbert SJ. Cognitive offloading. *Trends Cogn Sci.* 2016;20(9):676-88.
10. Bastani H, Bastani O, Sungu A, Ge H, Kabakçı Ö, Mariman R. Generative AI without guardrails can harm learning: evidence from high school mathematics. *Proc Natl Acad Sci U S A.* 2025;122(26):e2422633122.
11. Lee HP, Sarkar A, Tankelevitch L, Drosos I, Rintel S, Banks R, et al. The impact of generative AI on critical thinking: self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems.* New York: ACM; 2025.
12. Ahluwalia PS. Algorithmic bias and social stratification: how AI shapes inequality in digital societies. *Siddhanta's Int J Adv Res Arts Humanit.* 2025:28-37.
13. Ahluwalia PS. The role of artificial intelligence in combating plagiarism in scholarly publishing. *Shodh Prakashan: J Humanit Soc Sci.* 2025:26-30.
14. Bender EM, Gebru T, McMillan-Major A, Shmitchell S. On the dangers of stochastic parrots: can language models be too big? In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21).* New York: ACM; 2021. p. 610-23.
15. Shumailov I, Shumaylov Z, Zhao Y, Papernot N, Anderson R, Gal Y. AI models collapse when trained on recursively generated data. *Nature.* 2024;631(8022):755-9.
16. Mednick SA. The associative basis of the creative process. *Psychol Rev.* 1962;69(3):220-32.
17. Amabile TM. The social psychology of creativity: a componential conceptualization. *J Pers Soc Psychol.* 1983;45(2):357-76.
18. Swanson DR. Fish oil, Raynaud's syndrome, and undiscovered public knowledge. *Perspect Biol Med.* 1986;30(1):7-18.
19. Bjork EL, Bjork RA. Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning. In: *Gernsbacher MA, Pew RW, Hough LM, Pomerantz JR, editors. Psychology and the real world: essays illustrating fundamental contributions to society.* New York: Worth Publishers; 2011. p. 56-64.
20. Parasuraman R, Manzey DH. Complacency and bias in human use of automation: an attentional integration. *Hum Factors.* 2010;52(3):381-410.
21. Ahluwalia PS, Elliason EK. Decolonizing sociology: indigenous epistemologies and the rewriting of sociological theory. *Siddhanta's Int J Adv Res Arts Humanit.* 2025:125-136.
22. Kaur A, Ahluwalia PS. Jnana yoga and epistemology: ways of knowing in yogic philosophy. *Sanatanodaya.* 2025:166-183.
23. Ahluwalia PS. Digital dopamine: the neuropsychology of social media addiction among adolescents. *Siddhanta's Int J Adv Res Arts Humanit.* 2025:93-109.
24. Ahluwalia PS. Doomscrolling and mental fatigue: cognitive overload in the era of crisis media. *Siddhanta's Int J Adv Res Arts Humanit.* 2025:250-270.
25. Ahluwalia PS, Singh D. The rise of digital minimalism: a new form of social resistance? *Siddhanta's Int J Adv Res Arts Humanit.* 2025:218-239.

26. Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ*. 2021;372:n71.
27. Tricco AC, Lillie E, Zarin W, O'Brien KK, Colquhoun H, Levac D, et al. PRISMA extension for scoping reviews (PRISMA-ScR): checklist and explanation. *Ann Intern Med*. 2018;169(7):467-73.
28. Ahluwalia PS. Enhancing systematic transparency through PRISMA-ScR: a scoping review approach for emerging researchers. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:1-27.
29. Ahluwalia PS. Structured exploration through PRISMA-ScR: a transparent method for mapping research landscapes. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:137-167.
30. Ahluwalia PS. Plagiarism: detection techniques and tools. *Siddhanta's Int J Adv Res Arts Humanit*. 2024:1-11.
31. Ahluwalia PS. How we make plagiarism-free content: theory and practices. *Shodh Prakashan: J Humanit Soc Sci*. 2025:46-53.
32. Ahluwalia PS. Ethical issues in social science research. *Shodh Patra: Int J Multidiscip Stud*. 2024:1-8.
33. Ahluwalia PS, Khandey PD. From NFT to artifact: the changing value of art in a post-digital economy. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:147-156.
34. Ahluwalia PS, Singh D. Memes as cultural capital: humor, identity, and resistance in digital communities. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:115-124.
35. Ahluwalia PS. Influencer culture and neoliberal identity: a sociological critique of self-branding on social media. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:78-91.
36. Ahluwalia PS. The rise of techno-authoritarianism: digital surveillance and sovereignty in the Global South. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:34-54.
37. Ahluwalia PS, Kumar A. Platform poverty: the social consequences of algorithmic invisibility on gig workers. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:240-249.
38. Ahluwalia PS. The rise of quiet quitting: work ethic, alienation, and resistance in modern employment. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:92-106.
39. Ahluwalia PS. Artificial intelligence in healthcare: the role of artificial intelligence in diagnostic medicine. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:164-175.
40. Ahluwalia PS. Redesigning patient engagement in AI-driven healthcare systems: an analysis. *Siddhanta's Int J Adv Res Arts Humanit*. 2025:203-217.
41. Singh P. Integrating machine learning for smart facility management in healthcare hospitality environments. *Shodh Patra: Int J Multidiscip Stud*. 2025:82-92.
42. Ahluwalia PS. Artificial intelligence-driven prediction of antimicrobial resistance patterns: integrating machine learning with genomic epidemiology for precision antibiotic stewardship. *J Integr Res*. 2025:97-122.
43. Ahluwalia PS. Assessing the impact of digital transformation on hospital management: a systematic review. *Int J Healthc Hosp Manag Stud*. 2025:1-16.
44. Ahluwalia PS. Google programs for educational upliftment: a comprehensive analysis. *Shodh Prakashan: J Educ Res Train*. 2025:59-64.
45. Ahluwalia PS. Educational drawbacks of the New Education Policy–2020. *Shodh Prakashan: J Educ Res Train*. 2025:43-52.
46. Ahluwalia PS. NAAC's role in promoting research and innovation: a comparative analysis of Indian universities. *Siddhanta's Int J Adv Res Arts Humanit*. 2023:47-70.
47. Bala N, Ahluwalia PS. Neuroplasticity and yoga: how practice reshapes the brain. *Sanatanodaya*. 2025:286-301.