How AI Writing Tools Change Student Learning Habits: A Review of Previous Research

Vaishnavi Ghanwat[1], Sakshi Mishale[2], Priyanka Murme[3], Vaishnavi Khalate[4]
[1,2,3,4]*Computer Science, MIT Art's, Commerce, &Science College Alandi*
doi.org/10.64643/ijirt.208438-459

***Abstract*—Artificial Intelligence (AI) writing tools are increasingly becoming part of students' learning and academic writing practices. Previous research shows that these tools can help students generate ideas, improve language, revise their work, and receive quick feedback. Recent research on university students has also found that AI-supported writing can improve writing fluency, accuracy, and vocabulary. For example, a 2025 study of IT students learning Russian as a second language reported a 23% increase in essay length, a 31% reduction in language errors, and an improvement in lexical diversity after using AI writing tools with guided learning activities. The study also found that students considered AI useful for improving writing clarity and revision, while concerns about over-reliance and ethical use were also reported. Importantly, 70% of students showed greater awareness of authorship, transparency, and academic integrity. These findings suggest that AI writing tools can support learning when they are used with proper guidance rather than replacing students' own efforts. This review paper examines how AI writing tools are changing student learning habits, particularly in writing, information searching, revision, independent learning, and engagement. It also discusses the benefits, challenges, ethical concerns, and possible effects of excessive dependence on AI. Overall, previous research suggests that responsible and guided use of AI can support students' learning while maintaining critical thinking, creativity, and academic integrity.**

***Index Terms*—artificial intelligence, AI writing tools, student learning habits, academic writing, self-regulated learning, academic integrity, critical thinking, generative AI, ChatGPT, higher education**

I. INTRODUCTION

The rapid diffusion of generative artificial intelligence (AI) into everyday writing workflows has introduced one of the most significant disruptions to student learning habits in recent decades. Tools originally designed for narrow tasks such as grammar correction have evolved into general-purpose writing assistants capable of drafting, restructuring, summarizing, and even generating entire essays from short prompts [1], [4]. This evolution has occurred so quickly that instructional practice, assessment design, and institutional policy have struggled to keep pace [13].

Understanding the effect of these tools on student learning habits is not simply a question of whether AI improves or harms writing quality. It is a question of behavioural change: how students plan, draft, revise, seek feedback, and manage academic honesty once an always-available AI collaborator is part of the process [3], [11]. Early research treated AI writing tools primarily as grammar checkers with clear, bounded benefits [8]; more recent research treats them as general cognitive collaborators whose influence extends into critical thinking, motivation, and academic integrity [9], [16].

The stakes of this question extend beyond individual assignments. Writing instruction has long been treated by educators not merely as a mechanism for producing a finished document, but as a vehicle for developing analytical reasoning, argument construction, and independent thought [9]. If AI tools alter the habits through which students practice these underlying skills — for example, by shifting effort away from generating original arguments and toward evaluating or lightly editing AI-generated arguments — then the consequences of AI adoption could extend well beyond writing courses into any discipline that relies on writing as a vehicle for thinking [14]. This is part of why faculty concern, discussed in Section V, is not limited to composition instructors but extends broadly across disciplines that assign written work.

This paper synthesizes previous research on this topic to answer three questions. First, what learning habits have been most clearly affected by AI writing tool adoption? Second, which effects are consistently positive, which are consistently negative, and which remain contested? Third, what factors explain why the same technology produces beneficial outcomes for some students and harmful ones for others? The remainder of the paper is organized as follows: Section II outlines the review methodology; Section III categorizes the tools under review; Sections IV and V examine positive and negative effects respectively; Section VI discusses moderating factors; Section VII presents a conceptual framework and supporting visual data; Section VIII discusses implications; Section IX addresses limitations; and Section X concludes.

II. REVIEW METHODOLOGY

This review follows a narrative synthesis approach rather than a formal systematic review protocol, given the breadth of disciplines (composition studies, educational psychology, computer science education, and higher-education policy) that have published relevant work since 2021. Sources were drawn from peer-reviewed journals, systematic reviews, preprint repositories, and large-scale institutional surveys published between 2021 and 2026 [1]– [18]. Priority was given to studies reporting empirical outcomes (pre/post writing scores, survey means, or behavioural measures) over purely theoretical commentary, and to more recent publications when findings on the same question diverged over time.

Sources were grouped thematically into three categories: (1) studies documenting positive effects on writing performance and engagement; (2) studies documenting negative effects related to over-reliance, critical thinking, and academic integrity; and (3) studies examining moderating variables such as self-regulated learning, AI literacy, and institutional policy. This thematic structure forms the basis for Sections IV through VI.

III. OVERVIEW OF AI WRITING TOOLS IN EDUCATIONAL USE

The tools discussed across the reviewed literature fall into three broad functional categories. The first is correction-oriented tools such as Grammarly, which provide sentence-level feedback on grammar, style, and clarity [1], [8]. The second is generative tools such as ChatGPT and Gemini, which can produce original drafts, outlines, or paraphrased content from a prompt [2], [5]. The third is hybrid platforms, such as Scholar's AI-assisted peer-review environment, which combine generative feedback with structured peer- and instructor-review workflows [1]. Reviewed studies show adoption spanning all three categories, often simultaneously, with students frequently combining a generative tool for idea development with a correction tool for final proofreading [4], [9].

The distinction between these categories matters for the review that follows because the reviewed studies do not treat them as interchangeable. Correction-oriented tools are described almost uniformly favourably across the literature, largely because their scope is narrow and their output is easy for a student to evaluate against their own intent a suggested grammar fix is either correct or not, and accepting it does not require ceding authorship of the underlying idea [8]. Generative tools, by contrast, are described with far more ambivalence, precisely because their output can substitute for, rather than merely refine, a student's own reasoning [2], [9].

Hybrid platforms occupy a middle position: because they embed generative AI feedback within a structured review workflow that also includes human peer or instructor input, several studies suggest that the surrounding pedagogical structure not the AI component alone determines whether the tool supports or substitutes for learning [1]. This distinction between tool category and pedagogical context reappears throughout Sections IV through VI and is a central input to the conceptual framework developed in Section VII.

Comparison of AI Writing Tool Categories

Table I summarizes the three functional categories identified across the reviewed literature, the primary mechanism by which each is reported to influence student writing habits, and representative sources.

TABLE I. AI Writing Tool Categories and Mechanisms of Influence

Category	Representative Tools	Primary Mechanism	Key Sources
Correction-oriented	Grammarly	Sentence-level grammar, style, clarity feedback	[1], [8]
Generative	ChatGPT, Gemini, QuillBot	Drafting, outlining, paraphrasing, idea generation	[2], [5]

Hybrid feedback platforms	Scholar AI Helper	AI-augmented peer and instructor review workflows	[1]

IV. POSITIVE EFFECTS ON STUDENT LEARNING HABITS

The most consistent finding across the reviewed literature is that AI writing tools measurably improve task fulfilment and organizational quality in student writing. A content-analysis study using pre- and post-tests found significant score gains in task fulfilment and organization following AI-assisted writing instruction [2]. Similarly, a qualitative investigation of AI-supported peer feedback reported that students who engaged with AI-enhanced feedback improved their final assignment grades by an average of fifteen percent compared to prior AI-free submissions, alongside greater confidence in academic communication [5].

Beyond scores, several studies report a change in the underlying habit of revision itself. Rather than treating a first draft as close to final, students using AI feedback tools engage in more iterative, dialogic revision, returning to a draft multiple times in response to targeted, immediate feedback [1]. This dialogic quality is significant because immediate feedback loops are historically difficult for instructors to provide at scale; AI tools appear to partially fill this gap by creating a sense of ongoing feedback rather than a single graded checkpoint [1].

A related and frequently cited benefit is support for higher-order writing skills. Research on paraphrasing and clarity found that AI-assisted feedback helped students move beyond surface corrections toward more substantive revisions of argument structure and coherence [1]. Other work emphasizes idea generation, noting that generative tools such as ChatGPT support brainstorming and creativity in the early stages of the writing process, which some students found reduced the anxiety associated with starting a writing task [2].

Taken together, these findings indicate that AI writing tools change learning habits in a broadly positive direction when the target is procedural: students draft earlier, revise more often, and report higher confidence. The open question, addressed in the next section, is whether these procedural gains come at the cost of independent cognitive effort.

Effects on Specific Learning Habits

Beyond aggregate score gains, the reviewed literature identifies four discrete habits that shift consistently once AI writing tools enter a student's workflow: planning, drafting, feedback-seeking, and time management.

Planning behaviour changes in that student’s report using generative tools to externalize early-stage brainstorming, producing rough outlines or bullet-point idea lists before committing to prose [2]. This appears to lower the psychological barrier to starting a writing task, an effect several studies link to reduced writing-related anxiety [2], [9]. Drafting behaviour shifts from a single linear pass toward a more iterative loop in which a student produces a partial draft, submits it for AI feedback, and revises before continuing, rather than completing an entire draft before seeking any feedback [1].

Feedback-seeking habits also change quantitatively: because AI feedback is available on demand rather than on an instructor's schedule, students in several studies sought feedback substantially more often per assignment than comparable cohorts without AI access, and reported that the immediacy of this feedback — rather than its content alone — was the primary driver of increased engagement [1], [5]. Finally, time-management habits shift in mixed ways; some students report that AI-assisted drafting frees time for deeper revision, while others report that easy access to AI-generated text encourages later starts on assignments, since the perceived time cost of producing a first draft is lower [9], [16].

V. NEGATIVE EFFECTS AND EMERGING CONCERNS

A. Over-Reliance and Critical Thinking

The clearest counter-pattern in the literature is over-reliance. Multiple studies warn that when AI tools are used to generate content wholesale rather than to support a student's own drafting process, students fail to practice the independent articulation of ideas that writing is meant to develop [9], [12]. A 2026 path-analysis study found that heavier AI tool use was associated with lower self-reported critical thinking unless moderated by strong self-regulated learning skills [11]. Large-scale survey data reinforce this concern at the perception level: in a global survey of students, seventy percent reported a perceived decline in their own critical thinking associated with AI tool use, even though fifty-five percent simultaneously reported cognitive advantages such as improved writing and problem-solving [18].

Faculty perceptions echo this tension even more strongly. A large 2026 College Board survey of postsecondary faculty found that eighty-eight percent were concerned about student overreliance on automation, and eighty-four percent agreed that AI reduces students' critical thinking, originality, and deep engagement with course material [13]. These are not fringe concerns: they represent a near-consensus view among instructors actively teaching writing-intensive courses.

B. Academic Integrity

Academic integrity is the second major negative theme. The same College Board survey found that ninety-two percent of faculty were concerned about plagiarism or dishonesty facilitated by AI [13]. Separately, industry survey data show a sharp rise in instructor-reported dishonesty: seventy-two percent of instructors reported believing some students had cheated in 2021, rising to ninety-six percent making the same claim by 2024 — a shift that closely tracks the mainstream availability of generative AI tools during the same period [15].

Student self-report data corroborate this trend from the other side: over half of surveyed students said they believed cheating had increased relative to the previous year, with nearly a quarter reporting a significant increase [15].

Importantly, students themselves are not uniformly unconcerned about this shift. One survey of English-majored students found that eighty percent worried that overuse of AI tools would erode their sense of independent thought, even though sixty-two percent admitted to using such tools regularly [18]. Another study found that sixty-one percent of respondents supported integrating AI into the learning process only under explicit limitations, suggesting that students themselves recognize the need for guardrails rather than unrestricted use [12].

C. Equity, Privacy, and Data Concerns

A smaller but recurring theme concerns equity and privacy. Because AI writing tools are trained on large, uncredited text corpora, some scholars raise concerns about biases embedded in AI systems that could disadvantage certain student populations if left unaddressed [1]. Others note that tools requiring account access or data submission raise privacy concerns, particularly when institutions have not established clear data-governance policies for student work processed by third-party AI systems [1], [16]. Survey participants in one Sub-Saharan African university sample expressed moderate but consistent concern about data privacy (mean 3.09 on a 5-point scale) and transparency in AI data processing (mean 3.30), indicating that these concerns are not limited to well-resourced institutional contexts [16].

D. Disciplinary and Demographic Variation

The magnitude of both positive and negative effects is not uniform across student populations. Studies of language-learning contexts report particularly strong reliance on generative tools for idea generation and cultural or idiomatic phrasing, alongside calls for AI systems with greater cultural and emotional sensitivity to better support second-language writers [5]. Studies situated in computing and engineering education emphasize a distinct risk: when AI tools are used not only for prose but also for generating technical explanations, literature summaries, and even code, the erosion of independent problem-solving ability may extend beyond writing into core disciplinary competencies [14]. Socioeconomic context also appears to matter; a study conducted across Sub-Saharan African universities found ethical-concern and privacy-concern scores broadly comparable to those reported in higher-resource contexts, suggesting that concerns about over-reliance and integrity are not confined to any single institutional setting but generalize across regions with very different levels of technological infrastructure [16].

VI. MODERATING FACTORS: WHY EFFECTS DIVERGE

Given that the same class of tools produces both the procedural gains described in Section IV and the cognitive and integrity concerns described in Section V, a central question for this review is what explains the divergence. Three moderating factors recur across the literature.

First, self-regulated learning (SRL) capacity consistently predicts whether AI use translates into learning gains or learning substitution. Students with stronger planning, monitoring, and evaluation habits tend to use AI tools as a supplement to their own reasoning, while students with weaker SRL are more likely to accept AI output uncritically [3], [11]. Second, AI literacy understanding what these tools can and cannot reliably do appears to buffer against both over-reliance and misplaced trust in AI-generated content [3].

Third, institutional policy and classroom norms shape usage patterns; students who receive explicit guidance on acceptable AI use report more deliberate, bounded use than students operating without clear guidelines, and this factor also moderates the ethical-concern outcomes reported in Section V [14], [16].

These three factors SRL/AI literacy, frequency and purpose of use, and institutional policy form the mediating layer of the conceptual framework presented in the following section.

Habit Formation Over Time

A further moderating consideration, less thoroughly studied but consistently raised, is the trajectory of habit formation over an academic term or program. Several authors distinguish between AI use as a scaffold heavily relied upon early in a course and gradually withdrawn as competence and confidence increase and AI use as a permanent crutch, in which reliance does not diminish over time and may deepen [3], [11].

A longitudinal case study of independent writing skill development found that students who used AI tools primarily as a scaffold in early coursework retained stronger independent revision skills over a full academic year than students who continued heavy AI use throughout, even after controlling for initial writing ability [6]. This distinction scaffold versus crutch offers a more actionable lens for instructors than the frequency of use alone, since two students who use AI tools equally often may be on very different developmental trajectories depending on whether that use is fading or persisting.

Measurement Approaches Used in Reviewed Studies

The strength of evidence behind the patterns described in Sections IV through VI depends partly on how each underlying study measured its outcomes, and it is useful to make these approaches explicit before turning to the synthesis in Section VII. Three broad measurement approaches recur. The first is pre/post writing-score comparison, in which the same cohort's writing is scored before and after a period of AI-assisted instruction using a standardized rubric; this approach produced the task-fulfilment and organization gains discussed in Section IV but is limited by the absence of a no-AI control group in several studies [2]. The second is large-sample perception surveys, typically using five-point Likert scales, which produced the percentage figures underlying Figs. 1 and 2 and the critical-thinking and integrity findings in Section V; these are efficient at scale but measure perceived rather than directly observed behaviour [13], [15], [16], [18]. The third is qualitative case-study and interview methods, which produced the richer process-level findings on drafting, revision, and feedback-seeking habits discussed under Section IV, at the cost of smaller sample sizes and reduced generalizability [1], [5].

Few reviewed studies combine all three approaches within a single design, and this methodological fragmentation is itself a relevant finding: it means that the positive, procedural findings in this review rest more heavily on direct performance measurement, while the negative, cognitive and integrity-related findings rest more heavily on self-report. This asymmetry does not invalidate the negative findings, since perception data on a phenomenon like academic dishonesty is often the only ethically and practically feasible measurement approach, but it does mean the two halves of the picture presented in this review are not measured on fully equivalent footing, a point revisited in Section IX.

VII. CONCEPTUAL FRAMEWORK AND SUPPORTING DATA

Fig. 3 synthesizes the findings of Sections IV through VI into a single conceptual framework. AI writing tool adoption does not act on learning outcomes directly; rather, its effect is routed through three mediating conditions the student's self-regulated learning and AI literacy, the frequency and purpose of their tool use, and the institutional policy environment in which that use occurs. Depending on how these mediators combine, the same underlying technology produces the positive outcomes documented in Section IV, the negative outcomes documented in Section V.A, or the integrity risks documented in Section V.B.

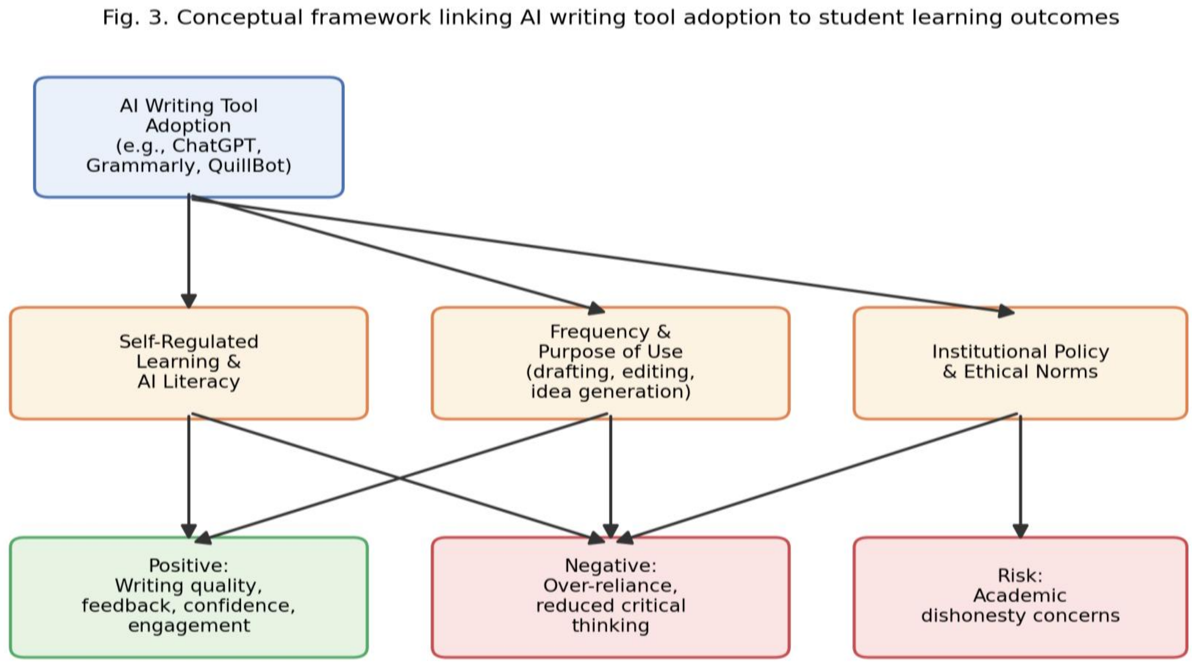

Fig. 3. Conceptual framework linking AI writing tool adoption to student learning outcomes.

This framework is consistent with the quantitative pattern visible in faculty and institutional survey data. Fig. 1 summarizes results from a large-scale 2026 faculty survey, showing that concern is highest for academic integrity (92%), followed by overreliance on automation (88%) and reduced critical thinking or originality (84%) [13].

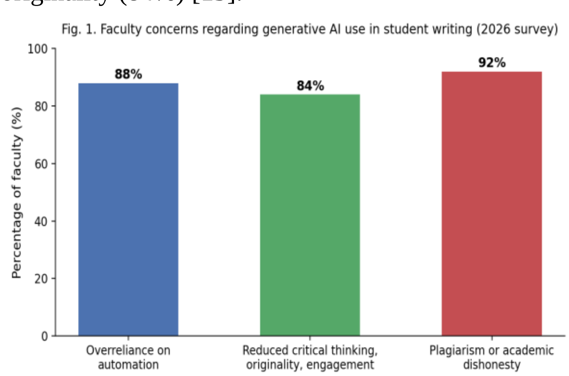

Fig. 1. Faculty concerns regarding generative AI use in student writing (2026 survey) [13].

The consistency of these figures all clustered in the mid-to-high 80s and low 90s indicates that faculty do not view these as three separate, independent worries but as facets of a single underlying concern about diminished independent effort.

Fig. 2 places this concern in temporal context using industry survey data comparing instructor-reported academic dishonesty in 2021, prior to widespread generative AI availability, against 2024 figures. The rise from seventy-two percent to ninety-six percent of instructors reporting at least some student cheating represents one of the more striking behavioural shifts in the reviewed literature, and it closely tracks the timeline of mainstream generative AI adoption in higher education [15].

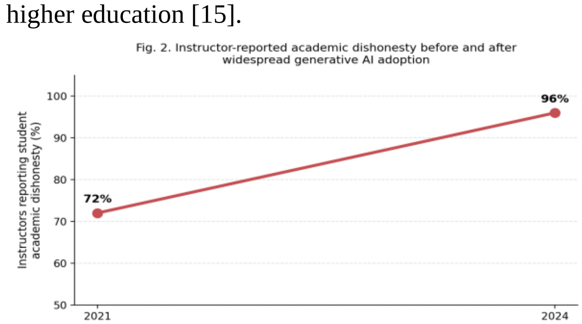

Fig. 2. Instructor-reported academic dishonesty before and after widespread generative AI adoption [15].

Faculty–Student Perception Gap

A notable cross-cutting theme in the literature is a persistent gap between how faculty and students perceive the same AI-driven changes in learning habits. Faculty concern, as shown in Fig. 1, clusters heavily around loss of independent effort, with nearly universal agreement on the risks of overreliance and academic dishonesty [13]. Students, by contrast, tend to report a more instrumental view: a majority describe AI tools as helpful for idea generation, time savings, and understanding difficult content, while a smaller but still substantial share separately acknowledge the same risks faculty raise [15], [18]. This is not a simple disagreement about facts — both groups often cite similar concerns when asked directly — but a difference in salience: faculty encounter the aggregate pattern across many students and assignments, while individual students evaluate their own use case, which they are more likely to view as measured and well-justified even when survey data show the same students holding generalized worries about their peers' behaviour [15]. This perception gap has practical implications for policy design, discussed further in Section VIII, because policies drafted primarily from faculty concern without corresponding student input risk being perceived as punitive rather than supportive.

VIII. IMPLICATIONS FOR PRACTICE

Beyond course-level design, the reviewed literature also points to implications for how writing itself is taught as a skill. If, as several studies suggest, the primary risk of AI adoption is substitution of independent idea generation rather than degradation of surface-level writing mechanics [9], [11], then instructional time may be better spent on argument construction, evidence evaluation, and structuring a position — the habits most exposed to substitution — rather than on grammar and mechanics, which AI tools already support well and which correction-oriented tools have supported uncontroversial for years [8]. This reallocation of instructional emphasis is a direct, practical consequence of the pattern identified in Sections IV and V, rather than a general recommendation about educational technology.

The reviewed evidence suggests that blanket bans and unrestricted permission are both poorly matched to the underlying pattern identified in this review. Because outcomes are mediated by SRL, purpose of use, and institutional policy rather than by the technology itself, interventions that target these mediators are likely to be more effective than interventions that target the tools directly [11], [14]. Concretely, the literature points toward three practical directions: (1) embedding explicit AI-literacy instruction alongside writing instruction, so students understand both the capabilities and limitations of these tools [3]; (2) redesigning assessment to reward process (drafts, reflection, revision history) rather than only final-product quality, which reduces the incentive for wholesale generation [9], [14]; and (3) establishing clear, course-level policies on acceptable AI use rather than relying on institution-wide rules alone, since student behaviour appears more responsive to specific, proximate guidance [16].

A fourth direction, less prominent in the literature but implied by the scaffold-versus-crutch distinction discussed in Section VI, is intentional fading: designing early assignments to permit or even encourage AI-assisted drafting, while progressively restricting AI involvement in later assignments so that independent competence is demonstrated by the end of a course or program [3], [6]. This mirrors long-established scaffolding practice in other skill domains, where a support is introduced deliberately and removed deliberately, rather than left in place indefinitely or withheld entirely, and several authors argue that writing instruction has not yet caught up to this model with respect to AI tools specifically [3]. This approach treats AI tools as training wheels rather than as either a permanent aid or a banned shortcut, and aligns assessment design with the observed developmental trajectory of habit formation rather than with a fixed, course-long policy.

These four directions are not mutually exclusive, and the reviewed literature suggests they are most effective in combination rather than in isolation. AI-literacy instruction without a corresponding change in assessment design, for example, may improve students' understanding of a tool's limitations without changing the incentive structure that makes wholesale generation attractive under time pressure [3], [9]. Conversely, process-oriented assessment without any AI-literacy component risks penalizing students for a lack of technical understanding of tools they were never taught to use critically. The faculty–student perception gap discussed above further suggests that whichever combination of these directions an institution adopts, the policy is more likely to be followed, rather than quietly circumvented, if students are involved in its development rather than presented with rules developed exclusively from instructor concerns [15], [18].

Institutional and Policy Responses

Institutional responses documented in the reviewed literature range considerably in scope and rigor. At one end, some institutions have issued broad, principle-level guidance encouraging instructors to set their own course-level AI policies, effectively delegating the mediating-policy factor identified in Section VI to individual faculty members [14].

At the other end, some computing and engineering programs have begun specifying tool-level and task-level permissions for example, permitting AI-assisted debugging while restricting AI-generated literature reviews reflecting an emerging recognition that a single institution-wide rule cannot adequately cover the range of tasks in which AI tools might be used [14].

A recurring recommendation across multiple sources is that policy should specify not only whether AI use is permitted, but for which stage of the writing process (planning, drafting, editing) it is permitted, since the reviewed evidence in Sections IV through VI suggests that the learning consequences of AI use differ substantially by stage rather than being uniform across the writing process as a whole [1], [9], [14].

A second policy theme concerns detection and enforcement. Several sources note that students are aware of, and sceptical of, the reliability of AI-content detection and plagiarism-similarity tools, with survey respondents expressing moderate concern about the accuracy of such detection systems [16]. This scepticism matters for policy design: enforcement mechanisms perceived as unreliable may undermine student trust in institutional AI policy more broadly, reinforcing the perception-gap dynamic discussed above rather than resolving it. Sources examining faculty responses to rising dishonesty rates (Section V.B) note a similar pattern, with the majority of instructors reporting that managing student AI use presents at least a minor ongoing challenge, indicating that enforcement infrastructure has not fully caught up with the scale of adoption reflected in Fig. 2 [13].

IX. LIMITATIONS OF THE REVIEWED RESEARCH

Several limitations recur across the primary studies synthesized here and should temper the strength of the conclusions drawn. Most studies rely on self-administered questionnaires, which are susceptible to social-desirability bias, particularly on sensitive questions about academic dishonesty [11]. Sample sizes vary widely across studies, from small qualitative case studies of a single course to surveys of several thousand participants, limiting direct comparability [4]. Additionally, because the underlying technology changes rapidly, findings from studies conducted even two to three years apart may reflect different generations of tools with materially different capabilities, meaning some earlier findings may not generalize to current AI systems [4], [6].

A further limitation concerns causal inference. The large percentage figures reported in Section V and visualized in Figs. 1 and 2 are correlational: they establish that faculty concern and reported dishonesty have risen alongside AI adoption, but the reviewed survey designs are not able to fully rule out other contributing factors over the same period, such as changes in assessment format, increased academic pressure, or shifts in reporting behaviour independent of actual behaviour change [13], [15]. Relatedly, this reviews itself relies on a narrative rather than a systematic synthesis methodology (Section II); while this allowed broader thematic coverage across disciplines and publication types, it does not provide the same protection against selection bias in source inclusion that a formal systematic review with pre-registered search criteria would offer. Readers should therefore treat the framework in Fig. 3 as an organizing synthesis of the current literature rather than as a validated causal model.

Finally, publication patterns themselves may introduce bias. Several of the most widely cited studies in this space originate from institutions and research teams already invested in either promoting or critiquing generative AI adoption in education, which may influence both study design and the framing of results [4], [6]. Cross-disciplinary replication for example, comparing findings from composition-studies journals against findings from computer-science-education venues on the same underlying question remains limited, and would strengthen confidence in the generality of the patterns summarized in Table II.

Summary of Reviewed Effects

TABLE II. Direction of Change in Student Learning Habits Associated with AI Writing Tool Adoption

Learning Habit	Direction of Change	Key Evidence
Drafting / starting a task	Earlier, more frequent	[1], [2], [9]
Revision	More iterative and frequent	[1], [5]
Feedback-seeking	More frequent, more immediate	[1], [5]
Independent idea generation	Reduced without SRL/AI literacy	[9], [11], [18]
Critical evaluation of sources/claims	Reduced (perceived)	[11], [13], [18]
Academic-integrity compliance	Weakened; rising dishonesty reports	[13], [15], [16]

X. CONCLUSION

Taken as a whole, the eighteen sources synthesized in this review span composition studies, educational psychology, computer-science education, and large-scale institutional survey research, and they converge on a picture that is more nuanced than either optimistic or pessimistic accounts of AI in education typically allow. The consistency of the procedural gains described in Section IV, replicated across case-study, survey, and pre/post-test designs, gives reasonable confidence that AI writing tools do change early-stage writing behaviour in a measurable and largely constructive way. The consistency of the integrity and critical-thinking concerns described in Section V, replicated across faculty and student samples on different continents, gives similar confidence that these risks are real rather than a moral panic confined to a single institutional context. What the literature does not yet offer is a validated method for predicting, at the level of an individual student or course, which of these two patterns will dominate which is precisely the gap the framework in Fig. 3 is intended to make explicit rather than to resolve.

This review set out to characterize how AI writing tools are changing student learning habits. The evidence indicates a genuinely mixed picture rather than a simple net-positive or net-negative effect. Procedural habits drafting earlier, revising more often, engaging with feedback more frequently shift in a positive direction with AI tool adoption. At the same time, independent cognitive habits original idea generation, self-directed critical evaluation, and adherence to academic-integrity norms are placed under measurable strain, particularly in the absence of strong self-regulated learning skills or clear institutional guidance.

The conceptual framework proposed here suggests that the productive path forward is not a debate over whether to permit AI writing tools, but a deliberate design effort around the mediating conditions AI literacy, purpose of use, and policy that determine which of these two patterns dominates for a given student in a given course. The faculty–student perception gap identified in Section VII further suggests that this design effort cannot be conducted from either group's perspective alone: policies developed without student input risk being seen as arbitrary, while course norms developed without faculty consensus risk being inconsistently enforced. Future research would benefit from longitudinal designs that track individual students' AI-assisted and independent writing competence across an entire program, rather than a single course or semester, in order to determine more precisely when AI-assisted scaffolding fades into independent competence and when it instead hardens into long-term dependence.

REFERENCES

[1] S. A. B. (Scholar Project Team), “The impact of AI-driven tools on student writing development: A case study from the Scholar AI Helper Project,” *arXiv preprint arXiv:2501.08473*, 2025.

[2] “The impact of AI-powered writing tools on students' writing performance: A content analysis and future prospects,” ResearchGate, 2025.

[3] “Exploring how AI literacy and self-regulated learning relate to student writing performance and well-being in generative AI-supported higher education,” PubMed Central, National Library of Medicine, 2025.

[4] "The impact of generative AI on academic reading and writing: A synthesis of recent evidence (2023–2025)," Frontiers in Education, 2025.

[5] P. Bhamare, S. Shevkari, and P. Mane, "How machine learning with Python and artificial intelligence benefits Formula 1 teams," International Journal of Innovative Research in Technology, vol. 12, no. 6, 2025.

[6] "Exploring artificial intelligence and culture: Methodology for a comparative study of AI's impact on norms, trust, and problem-solving across academic and business environments," arXiv preprint arXiv:2510.11530, 2025.

[7] H. Aljunied, "The impact of artificial intelligence tools on academic writing instruction in higher education: A systematic review," Arab World English Journal, Special Issue on ChatGPT, pp. 26–55, Apr. 2024.

[8] "The impact of AI-driven tools on student writing development," Online Journal of Communication and Media Technologies, 2025.

[9] "The role of AI in academic writing: Impacts on writing skills, critical thinking, and integrity in higher education," Education Sciences, vol. 15, no. 9, Art. no. 247, 2025.

[10] G. N. Taur, M. M. Vidhate, and S. S. Shevkari, "Deep learning meets biometrics: A comparative study of modern identification techniques," International Journal of Innovative Research in Technology, vol. 12, no. 2, pp. 545–552, 2025.

[11] "The use of AI tools in academic writing and its associations with college students' critical thinking and academic integrity: A path analysis involving self-regulated learning and innovative behaviours," Frontiers in Psychology, 2026.

[12] "The impact of AI writing tools on academic integrity: Unveiling English-majored students' perceptions and practical solutions," ResearchGate, 2025.

[13] College Board, "New College Board research: Faculty express near-universal concern that student AI use undermines original writing and critical thinking," College Board Newsroom, 2026.

[14] S. Devikar, M. Hinge, and S. S. Shevkari, "Human vs. machine: A review of AI's impact on language learning interaction or replacing human interaction," International Journal for Multidisciplinary Research, vol. 7, no. 3, 2025, doi: 10.36948/ijfmr. 2025.v07i03.49459.

[15] John Wiley & Sons, "AI has hurt academic integrity in college courses but can also enhance learning, say instructors, students," Wiley Newsroom, 2024.

[16] "The role of over-reliance on AI in the negative consequences of student learning: The moderating effects of ethical concerns and institutional policies," Cogent Education, Taylor & Francis, 2025.

[17] "LLMs and childhood safety: Identifying risks and proposing a protection framework for safe child-LLM interaction," arXiv preprint arXiv:2502.11242, 2025.

[18] "Student perceptions of AI-assisted writing and academic integrity: Ethical concerns, academic misconduct, and use of generative AI in higher education," *MDPI*, vol. 1, no. 1, Art. no. 2, 2025.