

Agentic Displacement in Second Language Writing: The Acquisitional Consequences of Delegating Writing to Autonomous AI Browsers

Hala Fawzi^{1*}

¹Mohammed Jaber Al Ansari College for Teachers, University of Bahrain, Sakhir, Kingdom of Bahrain

*Corresponding author's email: hawzi@uob.edu.bh

 <https://orcid.org/0000-0001-9379-3190>

 <https://doi.org/10.54855/ijte.26642>

® Copyright (c) 2026 Hala Fawzi

Received: 17/06/2026

Revision: 29/09/2026

Accepted: 01/10/2026

Online: 14/10/2026

ABSTRACT

Keywords: agentic displacement, second language writing, self-regulated learning, agentic AI browsers, second language acquisition

Autonomous AI browsers such as Perplexity Comet and Copilot Mode in Microsoft Edge can receive a writing assignment, draft a response, and submit it with little language produced by the learner. To explain what this means for second language (L2) writing, this conceptual paper introduces the term agentic displacement. Full delegation transfers planning, writing, and revision to the agent, removing opportunities for the learner to produce language and regulate the writing process. The paper argues that this loss may limit language development. Delegation is a continuum; the paper traces its extreme, full replacement, through production-based accounts of acquisition. No empirical data are reported; the paper offers a framework, five testable propositions, and three dimensions of teacher practice: recognizing displacement, auditing tasks for delegability, and using pedagogical judgment to recover the learning a task was meant to produce. The central risk is developmental, not only ethical or procedural.

Introduction

Generative AI has entered second language (L2) writing, as Hyland (2026) observes, “with extraordinary speed and with remarkably little agreement about what, precisely, it means for the field” (p. 1). Part of the disagreement is definitional. Discussions tend to treat AI as a single category, yet the tools available to learners differ in a way that matters more than the technology they share. Some tools provide feedback on the learner’s writing; others write on the learner’s behalf. The first step, therefore, is to distinguish tools that assist from tools that execute, because the developmental consequences differ.

The distinction is clearest in agentic AI browsers. Agentic AI systems plan, reason across multiple steps, use tools, and pursue goals with minimal human oversight, which separates them from generative AI that responds to prompts (Dwivedi et al., 2026). Agentic browsers apply this architecture to web-based task execution, part of what Wiederhold (2026) calls the

agentic web, in which autonomous systems become the web's primary users. A student who delegates the full assignment brief to Perplexity Comet, or to ChatGPT Atlas before its withdrawal, is no longer using AI as a writing-support tool; the student is issuing a task to an autonomous agent that researches the topic, produces a draft meeting the assignment criteria, and, where the submission mechanism permits, completes the submission. The product's own documentation confirms this: OpenAI's documentation for ChatGPT Atlas (see Section 2.1) describes agent mode as taking actions in the browser, navigating, clicking, and completing multi-step flows on sites where the user is already signed in, subject to approval checkpoints on sensitive actions (OpenAI, n.d.-b). In Perplexity Enterprise, organization administrators can configure Comet Assistant's browser-control permissions, including whether it may click, navigate webpages, and fill forms for organization members; these permissions may be controlled case by case and by domain (Perplexity, 2026). Bozkurt et al. (2026) call the result the "ghost student", a surrogate of language model and agentic browser that completes assessments while "rendering the actual learner's presence optional" (p. 1). Teaching staff across the California Community Colleges system report agentic tools logging into Canvas, completing quizzes and discussions, and submitting work that appears "mostly indistinguishable from human activity" (California Virtual Campus [CVC], 2025, p. 1; see also Roitman, 2025). No prevalence data or empirical studies of agentic-browser use in L2 writing yet exist. The capability of these tools is documented; the scale of their use is not. The paper is therefore offered as a forward-looking conceptual framework for a serious pedagogical risk, and it specifies what evidence would confirm or limit it.

Three areas of research are relevant to this problem, but none directly addresses L2 writing. Academic integrity research has examined assignment outsourcing, its prevalence, predictors, and relationship to assessment design (Bretag et al., 2019; Newton, 2018), but asks about the ownership of the submitted work, not what the learner fails to learn by outsourcing it. Research on cognitive offloading explains how external tools take over cognitive functions and at what cost (Gerlich, 2025; Risko & Gilbert, 2016), but examines single functions, not the complete task. Work on hybrid human-AI regulation studies how learners and AI systems can share the regulation of learning productively (Järvelä et al., 2023; Molenaar, 2022), but assumes that control is divided and has not addressed the case in which it is handed over. None of the three answers the acquisitional question this paper raises: what happens to L2 development when the writing act, from forethought through submission, is done by the agent?

Approach

This is a conceptual paper. In Jaakkola's (2020) terms, it is a theory-synthesis paper: it integrates previously separate literature streams to build a new construct and offer a new view of the phenomenon. No new data are collected. The construct was built through conceptual synthesis across four literatures, each covering one part of the problem: (1) research on agentic AI systems, which supplies the operational definition of the tool class (Dwivedi et al., 2026; Kostopoulos et al., 2025; Sapkota et al., 2026); (2) the academic integrity literature on contract cheating, which supplies the nearest behavioral precedent and its predictors (Bretag et al., 2019; Newton, 2018); (3) self-regulated learning (SRL) research in L2 writing, which supplies the account of what regulated writing develops (Teng & Zhang, 2016, 2020; Zimmerman, 2000); and (4) production-based second language acquisition (SLA) theory, which supplies the mechanisms through which that development occurs (Manchón, 2011; Schmidt, 1990; Swain, 1985, 1995). The analysis follows the standard sequence for construct development: definition, differentiation from neighboring constructs, specification of boundary conditions, and derivation of falsifiable propositions. Section 2 defines agentic

browsers and locates delegation on a continuum. Sections 3 and 4 present the conceptual analysis in place of empirical results, distinguishing throughout between claims supported by direct evidence, claims supported by analogy, and theoretical predictions. The five propositions in Section 5 make the framework testable, and Section 6's implications for teacher knowledge and practice serve as the paper's discussion.

Agentic AI Browsers and the Delegation Continuum

What agentic browsers are

Agentic AI is the term now used for systems that are given a goal, break it into steps, and carry those steps out using tools with limited human direction (OECD, 2026; Sapkota et al., 2026). An agentic browser is the browser-based form of such a system. The browser no longer only displays pages for the user to act on. With a large language model built into it, the browser carries out multi-step tasks itself, such as searching, moving between sites, and entering information (Bozkurt et al., 2026; Garcia et al., 2026). Perplexity Comet and Copilot Mode in Microsoft Edge are current examples (Microsoft, n.d.; Perplexity, 2026). ChatGPT Atlas was a third until OpenAI discontinued it in August 2026 and moved its agent capabilities into ChatGPT and Codex (OpenAI, n.d.-a). Education research has so far studied agents in tutoring and other instructional roles (Kostopoulos et al., 2025), as workflow components (Dai et al., 2025), and as partners in learners' self-regulation (Salah et al., 2026). Salah et al. (2026) come closest to the present concern by identifying regulatory displacement, which occurs when AI does the regulating that learners need to practice themselves. This paper extends this concern to L2 writing and asks what happens when the displaced activity is language production itself.

Taxonomies of AI in education place systems on a spectrum of autonomy. Kostopoulos et al. (2025) order agents from reactive, through adaptive and proactive, to collaborative agents that plan and execute jointly with humans. The spectrum is instructive for what it does not contain. Even its most autonomous level keeps a human inside the task. The position taken here is that assignment-completing browser use is not a further point on this spectrum but a departure from it, because the property that orders the taxonomy, the depth of human participation in execution, drops to zero.

Current products qualify this autonomy with confirmation checkpoints, and that qualification needs an accurate description. Comet pauses at certain actions, as Atlas did, and a learner sometimes has to approve a step, which in principle keeps a human in the loop. In practice, a learner who wants a finished assignment approves almost automatically, and the checkpoint adds a click without adding real cognition. The claim defended in this paper is therefore not that learners press no keys. The claim is that the keys they press do not involve producing language.

Delegation as a continuum

AI use in writing cannot be looked at as a matter of use and non-use. A grammar checker used after a self-written draft takes over one editing function and displaces almost nothing. A chatbot asked to generate a paragraph that the learner evaluates, revises, and integrates displaces part of the performance phase but leaves forethought and self-reflection. An agentic browser given the assignment displaces everything between receiving the task and submitting it. These examples sit on a continuum of delegation, and it is this continuum, not any single position on it, that matters for analysis. Table 1 maps the main positions onto the SRL phases, the parts of the writing act that remain with the learner, and the predicted cost for acquisition.

Nothing in this paper claims that AI presence in writing is inherently damaging, only that the developmental consequences scale with how much of the writing act the learner gives away.

Existing studies do not present this continuum directly, but they provide the basis for its lower and intermediate positions. Warschauer et al. (2023) argue that the value of AI-generated text depends on what learners do with it, while Feng et al. (2025) locate reported benefits in planning, feedback and revision, where learners remain involved in composing. Table 1 reorganizes these uses according to what part of the writing act remains with the learner and which of Zimmerman's (2000) SRL phases are still engaged. Salah et al. (2026) treat agent autonomy as something to be bounded by the scope, timing and duration of intervention. The continuum developed here applies this graduated view to L2 writing and takes it to the limiting case of full delegation. The final column is this paper's own contribution: the acquisitional cost is a prediction derived in Section 4 and not a finding taken from the literature.

Table 1.

The delegation continuum: SRL phases engaged and predicted acquisitional cost

Position on the continuum	Writing act retained by the learner	SRL phases engaged	Predicted acquisitional cost
Learner writes; tool checks grammar and surface accuracy	Nearly all	Forethought, performance, self-reflection	Minimal
Learner writes; chatbot or tutor gives feedback on the draft (e.g., Feng et al., 2025)	All composing; evaluation shared	All three; reflection scaffolded	Low; feedback may support noticing
Learner plans and outlines; agent drafts; learner revises (Warschauer et al., 2023)	Planning and revision	Forethought and self-reflection; performance largely absent	Moderate; production functions engaged only in revision
Agent drafts from the brief; learner revises or post-edits (cf. Lee, 2020)	Revision only	Partial performance and self-reflection; forethought absent	High; see Section 3.3
Full delegation to an agentic browser, through to submission (Bozkurt et al., 2026)	None beyond the prompt	Task-completion regulation only; writing-oriented SRL absent	Maximal; the limiting case this paper analyses

Agentic Displacement

The construct

Agentic displacement is defined here as the pre-emptive removal of writing-oriented self-regulation from a writing task. It occurs when a learner delegates the task to an autonomous agent before any writing-oriented cognitive engagement with it begins. The definition presupposes an account of what is removed. Zimmerman (2000) models self-regulation as three cyclical phases: forethought, in which the learner analyzes task demands, sets goals, and

activates prior knowledge; performance, in which the learner executes the task while monitoring output; and self-reflection, in which the learner evaluates the outcome and recalibrates. Writing is a primary site for this model because all three phases are observable and teachable, and skilled writing is itself an exercise in self-regulation (Zimmerman & Risemberg, 1997).

In L2 writing specifically, the SRL apparatus is well documented. Teng and Zhang (2016) validated a multidimensional model of SRL writing strategies with 790 Chinese university learners, and five months of SRL-based writing instruction improved writing proficiency, with gains still present at the delayed post-test (Teng & Zhang, 2020). EFL learners in Vietnam, for example, report using SRL writing strategies alongside teacher feedback (Vu et al., 2022). Diary evidence from EFL writers records learners setting linguistic goals, activating vocabulary, monitoring accuracy, and reflecting on coherence across the writing cycle (Zhang et al., 2022). The collective picture is of writing development as regulated production: the learner's own goal-directed engagement with the task is the vehicle through which competence forms.

Tracing full delegation through the three phases shows what displacement removes. In the forethought phase, the agent interprets the task and plans the execution; the learner sets no linguistic goals and activates no prior knowledge. In the performance phase, the agent produces text; the learner monitors nothing, because there is nothing of theirs to monitor. In the self-reflection phase, the learner has no basis for evaluating decisions that were never theirs. The assignment is complete, but the cycle has never run.

What agentic displacement is not

Agentic displacement is useful as a construct only if it can be distinguished from three neighboring concepts: cognitive offloading, contract cheating, and hybrid human-AI regulation.

Cognitive offloading is the most clearly distinct. Offloading, in the canonical account, is the use of physical action or external resources to reduce the information-processing demands of a task (Risko & Gilbert, 2016). The defining feature is that the person retains direction of the task while a function is externalized. Recent empirical work in AI contexts suggests that heavy cognitive offloading is associated with weaker critical thinking (Gerlich, 2025), and that trust in AI predicts thoughtless use unless it is offset by critical thinking (Hou et al., 2025). Gerlich's (2025) evidence is correlational. A study of 666 participants found a negative correlation between frequent AI tool use and critical-thinking performance, mediated by offloading, but it relies partly on self-report and examines functions rather than whole tasks. On the account developed here, agentic displacement differs in scope and in the direction of control. The learner hands over the task itself, not one function within it.

Contract cheating serves as the most direct precedent here, because it is the one established practice in which a learner submits a complete text that another agent produced, so the developmental question it raises is the same one that agentic browsers raise (Bretag et al., 2019; Newton, 2018). Ghost-written assignments bypass planning, execution, and reflection entirely, so agentic displacement does not claim that the loss is new. A construct that ignored this precedent would claim a novelty it does not have. What is new is the way the task can now be handed over: quickly, privately, and inside the learner's own browser (Bozkurt et al., 2026; Roitman, 2025). Commercial contract cheating requires a second human, money, lead time, and exposure to a stranger. Each is a source of friction, and this paper treats that friction

as one likely reason why the practice remained a minority one, a hypothesis rather than an established finding.

Self-reports suggest a historic prevalence of nearly 3.5 percent of students, rising in recent samples (Newton, 2018), and the largest survey to date identified dissatisfaction with the teaching environment, perceived opportunity, and speaking a language other than English at home as the strongest correlates of outsourcing (Bretag et al., 2019). What agentic browsers change is not the mechanism of the loss but the economics of the practice, and this is the paper's novelty claim: agentic displacement does not name a new developmental loss; it names a new technical and institutional condition, automated, immediate, private, low-friction, and embedded in the learner's own browser, under which an old loss can occur at scale. Bretag et al.'s (2019) finding on language background also gives the L2 field a particular stake, since the learners most likely to outsource are, on the account developed in Section 4, those for whom the predicted acquisitional cost is highest.

Hybrid human-AI regulation describes the same variable from the opposite end. Molenaar (2022) observes that even sophisticated adaptive technologies tend to take over regulation from learners. She proposes designs in which control transfers gradually toward the learner as self-regulatory skill develops; Järvelä et al. (2023) extend the logic to shared regulation in collaborative learning. Displacement names the case in which that division collapses, and the transfer of control is total.

Boundary conditions and moderators

A construct is useful only when its boundaries are clear, and five limitations require more than brief acknowledgment. The first is hybrid use: the learner delegates the draft and then revises the agent's text. Revision reinstates parts of the performance and self-reflection phases and is acquisitionally preferable to untouched submission. Post-editing fluent text engages different linguistic operations from composing under communicative pressure, as research on machine-translation post-editing in L2 writing shows (Lee, 2020). The learner selects among, and repairs, formulations an external system has already settled, engaging the interlanguage only where a problem is noticed, so the retrieval and hypothesis-testing functions of production are weakly engaged. Comparing one's own attempt with a more proficient version can itself trigger noticing, as reformulation-based pedagogy shows (Swain & Lapkin, 2002), but reformulation works because the learner first writes and then compares. Delegate-then-edit removes the first term. Instructional design can restore it. A sequence in which the learner writes first and then compares the attempt with an agent-produced version converts post-editing into a designed noticing task. The boundary the construct draws is therefore between designed comparison, which preserves the learner's attempt, and spontaneous delegation, which removes it. No study has yet compared agentic post-editing with self-directed composition for acquisition. Hybrid use is placed at the high-cost end of the continuum in Table 1, with the prediction, not the assumption, that its cost falls between assistance and full delegation.

The second is prompt-writing. A careful prompt performs a residue of task analysis that may carry some metacognitive value, even in the first language (L1), but the residue is thin on the linguistic side. Prompts are short and formulaic, and nothing in them requires producing the structures the assignment was designed to elicit.

The third is the regulation a learner exercises over the agent itself. A learner who refines the goal, monitors the agent's progress, and evaluates the product is planning, monitoring, and evaluating. The construct asks what that activity regulates. Zimmerman's (2000) phases

develop competence in the domain being regulated, and in supervised delegation the regulated domain is task completion, not language production. The learner evaluates the agent's output as a reader, asking whether the text answers the brief, not whether their own language can express the intended meaning, and only the second problem bears on acquisition. Agent-directed regulation may still build a useful skill in managing AI tools, which fits with work on hybrid regulation (Molenaar, 2022). The claim is therefore limited. Supervision regulates task completion rather than language production, and that dissociation is stated as Proposition 5.

The fourth is proficiency, which connects to the distinction between acquisition and consolidation. For a novice, a writing task builds competence that does not yet exist; for an advanced learner, some tasks consolidate competence already formed. Displacement plausibly costs the novice more, and an advanced professional user of the L2 may approach delegation closer to the L1 deskilling case (Hyland, 2026). The construct's strongest claims therefore apply to learners for whom writing remains a primary engine of acquisition, and Section 5 treats proficiency as a moderator.

The fifth is context. L2 learners are not a homogeneous population. Learners in EFL settings with limited authentic exposure outside the classroom depend more heavily on classroom production opportunities than ESL learners surrounded by input, so the cost of forfeiting those opportunities is plausibly higher in EFL contexts. Task stakes matter: high-stakes, product-only assessment maximizes the incentive to delegate (Bretag et al., 2019). A related individual-difference point: for some learners, delegation may be a symptom of weak self-regulation rather than its cause, since SRL strategy use in L2 writing varies with motivational beliefs and self-efficacy (Teng, 2024), and the direction of effect between SRL capacity and delegation behavior is a question for the propositions below.

Acquisitional Consequences

Production-based acquisition mechanisms

For L2 learners, writing is more than a skill to be learned; it is a site where language itself is learned. Manchón (2011) terms this the writing-to-learn dimension: engagement in written production contributes to L2 development beyond the development of writing ability. An L1 professional who delegates writing loses practice. An L2 learner who delegates before competence has formed does not lose a skill but fails to build one. The mechanisms of that engine are production-based. Swain (1985) argued that comprehensible input alone cannot account for L2 development and that learners must also produce language under communicative pressure; Swain (1995) specified three functions of such output: noticing the gap between what one intends and what one can produce, testing hypotheses about form-function mappings, and reflecting metalinguistically on one's own production. Each function is triggered by the act of producing. An agentic browser produces text without the learner performing that act; the functions are not weakened but untriggered. Could a learner learn incidentally by reading the agent's polished output? Reading is input, and input has value, but the Output Hypothesis exists precisely because input alone under-delivers for productive competence. Agent-produced text returns the learner to that input-only condition. Recent work on AI-generated text in L2 writing reaches a compatible conclusion from the other direction. Warschauer et al. (2023) identify both affordances and contradictions in AI text for L2 writers and argue that its value depends on learners' active engagement with the text rather than passive acceptance of it, and Feng et al.'s (2025) review locates the documented benefits in feedback, planning, and revision rather than in text produced by the tool.

Schmidt (1990) locates the conversion of input to intake in attention to form. A noticing event occurs when a learner attempts a structure, perceives the mismatch between intention and capacity, and attends to the gap. The hypothesis is contested. Truscott (1998) argues that the evidence supports a far weaker role for conscious noticing than Schmidt claimed. The debate does not weaken the present argument, because displacement forecloses both readings: strong and weak formulations alike require the learner to engage with linguistic form during production, and under full delegation there is no production, no engagement, and nothing for attention, conscious or otherwise, to operate on.

The argument does not depend on a strong version of noticing. Usage-based, exemplar-based, and implicit-learning accounts would weigh the cost differently but not eliminate it. On those accounts, learning still depends on the learner's engagement with language through exposure, retrieval, and form-meaning matching, and full delegation cuts that engagement off. The paper does not argue that output is the only pathway to acquisition, only that a task designed to make the learner produce language loses its purpose when the learner produces none.

Effort, performance, and durable learning

Cognitive load theory distinguishes task demands from the processing that supports schema acquisition (Sweller, 1988). Applied to L2 writing, this distinction suggests that a well-designed task should direct learners' effort towards the language they are learning. Such a task makes the learner plan, find words, build sentences, and organize discourse at the same time. R. A. Bjork (1994) describes "desirable difficulties", conditions such as retrieval practice and spacing that can hinder immediate performance while supporting longer-term retention. Writing in an L2 is naturally one of those demanding conditions. Agentic browsers do not simply reduce the task load. They relocate the writing task, leaving little or no writing-related effort for the learner to perform.

The difference between performance and learning is important here, as they are not the same. The gap between them is where the cost shows. An L2 learner may perform well and produce a good text at a given time, but this does not mean that durable learning has taken place. Learning is durable and transferable change. E. L. Bjork and R. A. Bjork (2011) make this distinction clear by stating that conditions that improve immediate performance do not always support long-term learning. Fan et al. (2025) provide direct evidence of this divergence in AI-supported writing, in a randomized experiment with 117 university students. Learners supported by ChatGPT improved their essay scores more than learners supported by a human expert or by analytics tools, but they showed no corresponding advantage in knowledge gain or transfer. Fan et al. (2025) attribute this pattern to what they call "metacognitive laziness" (p. 489), a reliance on the tool that reduces learners' own regulatory engagement. The study measured knowledge of the essay topic rather than language development, so it supports the mechanism, not an acquisitional cost. And it examines assistance, in which learners remained responsible for producing the essay, not full delegation. This paper treats the extension to full delegation by agentic browsers as a prediction, not yet a finding. If assistance can improve scores without a matching gain in knowledge, full delegation may make this gap wider because the learner no longer carries the main writing work.

The incentive structure makes the predicted decoupling matter. Where the finished product is graded, and learning is only inferred, the rational strategy is to optimize the graded outcome, and productive struggle looks unnecessary even though it is the part of the task that supports development. An agentic browser meets the external requirements of a writing task at almost no cognitive cost. Where that option carries no visible penalty, nothing makes the effortful struggle feel worthwhile.

The scaffolding-execution distinction

This paper's argument is not against AI use in writing. Reviews of AI in L2 writing link its benefits to cases where the learner stays the writer, and the tool supports planning, revision, or feedback (Feng et al., 2025). Scaffolding keeps the learner inside the SRL cycle, while execution moves the task outside the learner's work. The variable that decides between them is where the learner stands in the task and what role the learner still has in the writing process, not how advanced the agentic browser system is.

Genre and task type as a moderator

The acquisitional cost of delegation is not uniform across writing tasks. The developmental cost depends on what the task expects the learner to produce. Some tasks are designed mainly to practice language, while others may ask students to communicate ideas, summarise information, or build an argument. This difference matters here. A highly controlled task that focuses on a specific grammar rule, such as a tense form, or on practicing a single structural move is designed mainly for linguistic practice. In this case, delegation removes the exact practice the task was designed to create. The developmental loss here is therefore high. In a more open content-driven assignment where the linguistic demands are spread out, such as a content-based essay or research synthesis, the loss is still present, but it is less concentrated. The learner still misses writing practice, but much of the actual difficulty lies in organizing ideas rather than forming specific sentences. This gradient follows from Manchón's (2011) writing-to-learn distinction. The more a task relies on the learner producing language, the more delegation takes away. Genre therefore belongs in the framework as a moderator alongside proficiency and context, and the propositions in Section 5 can be tested within genre to establish where the cost concentrates rather than assuming it falls evenly.

Propositions for Empirical Research

The construct is conceptual, but its claims are stated so that they can be tested. No studies of agentic browser delegation in L2 writing yet exist, so the most useful contribution a conceptual paper can make is to state its empirical commitments precisely enough for others to test them. Five propositions do so.

Proposition 1 (acquisitional cost of full delegation). Learners who fully delegate a writing task will show no pre-post gain on the linguistic structures the task targets, relative to proficiency-matched learners who complete the task with or without scaffolding AI. Full delegation can be operationalized through process evidence (no drafting activity in the document or LMS logs combined with submitted text beyond the learner's established profile or verified self-report). Development can be measured with pre- and post-task accuracy on targeted structures and standard complexity-accuracy-fluency indices, controlling for initial proficiency.

Proposition 2 (the gradient). Acquisitional cost will scale with delegation completeness across the positions in Table 1. Delegation completeness can be operationalized as the proportion of the submitted text composed by the agent, recoverable from process logs or keystroke data. The prediction extends the performance-learning divergence observed under assistance (Fan et al., 2025) and would be disconfirmed if delegate-then-revise use produced gains equivalent to learner-directed writing.

The operationalizations in Propositions 1 and 2 need caution, because process logs and keystroke data can misclassify behavior. For example, a learner might draft outside the monitored environment, use speech-to-text, paste from legitimate notes, or revise the agent's text in ways that hide where it came from. As a result, delegation should be triangulated rather than inferred from a single trace. Workable designs would combine version histories, short retrospective interviews or stimulated recall, screen-recorded writing in controlled studies, and comparisons with learners' established writing profiles. In privacy-constrained settings, researchers should prioritize learner consent and task-level evidence over covert monitoring.

Proposition 3 (prevalence and predictors). Full delegation through agentic browsers will occur at higher prevalence than commercial contract cheating. To test this prediction, measure prevalence using the same self-report instruments and incentivized-honesty methods used in the integrity literature. It should share the same predictors as contract cheating, such as perceived opportunity and dissatisfaction with the teaching environment (Bretag et al., 2019). The friction account in Section 3.2 depends on this comparison.

Proposition 4 (task design). Writing tasks that embed process evidence, such as documented drafting stages, in-class composition, and spoken justification of linguistic choices, will show lower full-delegation rates than product-only tasks of similar difficulty. A quasi-experimental design comparing matched course sections with and without process requirements, with delegation rate as the outcome, would test the claim directly.

Proposition 5 (relocated regulation). Learners who regularly delegate will show measurable gains in agent-management skill, operationalized as prompt specificity, error detection in agent output, and workflow evaluation, without corresponding gains on the interlanguage measures in Proposition 1. Confirmation would dissociate agent regulation from language learning regulation. Disconfirmation, gains on both, would show that supervised delegation carries acquisitional value the present framework denies it, and would require the construct to be revised.

Implications for Teacher Knowledge

Why teacher knowledge is the critical variable

If the account above is correct, the pressing pedagogical question is not how to detect agentic AI use. Detection tools assess statistical properties of text; they cannot observe whether the cognitive process that produces development occurred, and text produced by an agentic browser may be better formed, at the feature level, than the learner could produce. Kofinas et al. (2025) found in experiments at two universities that experienced markers could not reliably distinguish assessments with generative AI input from those without, recommending a shift from output-based to process-based and performative assessment. A systematic review of twenty studies reaches the same conclusion that distinguishing AI-generated from student-written text is unreliable (Dau & Lam, 2026). The productive question is what teachers must know and be able to do to protect the conditions that writing instruction exists to create. Teachers are the decision-makers closest to the point of loss: they design the tasks, observe the process or its absence, and decide how to respond. Each of the three dimensions that follow answers one face of the construct: functional recognition identifies displacement where it has occurred, critical audit reduces the delegability that invites it, and pedagogical judgment restores the acquisitional opportunity it removed. The three dimensions are derived from the theoretical analysis, not from evidence about what teachers currently know; whether

practicing L2 writing teachers possess them, and whether developing them reduces delegation or restores learning, are empirical questions on the same footing as the propositions above.

Functional recognition

The first dimension is the capacity to recognize signs that a submission was produced by an autonomous agent rather than by the student who submitted it. Textual evidence: Jiang and Hyland (2025) found that ChatGPT-generated essays deploy engagement markers in systematically different patterns from student essays, with markedly fewer questions, personal asides, and reader mentions. Pedagogical evidence: a teacher who knows a student's established interlanguage will notice a submission operating at a register and syntactic complexity the student has never produced in any other context, and a submission without drafts or revision traces, from a student who cannot account for its linguistic choices, presents a distinct evidential profile. Recognition of this kind is teachable through structured comparison of agent-produced and learner-produced texts.

Critical audit

The second dimension treats delegability as a property of task design that teachers can assess and reduce. A task consisting of a prompt, a word count, and an electronic submission mechanism can be assigned to an agentic browser in under a minute. Tasks with embedded process requirements resist. The integrity literature supports the design lever, since the assessment types students perceive as least likely to be outsourced are also the least frequently set (Bretag et al., 2019). Auditing tasks for delegability is not a surveillance exercise. Designs that resist delegation, including process portfolios, staged drafts, and oral defense, are recommended by SRL-informed writing instruction on independent grounds, and their effects on writing development are documented (Teng & Zhang, 2020).

A practical audit can come down to a few questions a teacher may ask before setting any writing task. Can an agent finish the task from the brief alone, leaving no trace of process? If so, the task is highly delegable. Does any part of the grade depend on evidence an agent cannot produce after the fact, such as in-class writing, dated drafts, or a learner's own account of their choices? If not, the task gives the teacher little evidence of how the writing was produced. If the learner delegates the task, what developmental opportunity is lost, and can the course recover it? Those three questions track exposure, detectability, and cost. Redesign first the tasks that score poorly on all three. The heuristic is a teaching tool, not a validated instrument, and whether its use lowers delegation rates is an empirical question of the kind Proposition 4 frames.

Process evidence such as draft histories, oral defense, and in-class writing protects acquisition, but it can disadvantage some learners, especially those with disabilities, high anxiety, limited or shared devices, unstable connectivity, or legitimate reliance on assistive technology. The response to agentic displacement should not rely on surveillance-heavy assessments. Learners should know what process evidence is expected, why it is expected, and what alternatives exist. To add flexibility, a process portfolio might accept typed drafts, audio reflections, a short teacher conference, or annotated revisions, provided that each still shows the learner engaging in the writing act.

Pedagogical judgment

The third dimension of this framework begins precisely where institutional policy ends. Once an instructor confirms a student delegated a task, the institutional response is usually straightforward. The harder pedagogical question is what the learner failed to develop and

how to give that back. This requires deep pedagogical judgment. A teacher needs to know the exact purpose of the original assignment. Was it targeting specific syntax? Was it building genre awareness or expanding vocabulary? Once that purpose is clear, the teacher must design a recovery task that recreates those conditions and makes the developmental cost of delegating clear to the learner, not only the integrity cost. If a learner only fears getting caught, that deterrent vanishes the moment AI detection fails, and current evidence indicates that it often does (Dau & Lam, 2026; Kofinas et al., 2025). A learner who genuinely understands that delegation undermines their own language development has a real reason to resist the technology. For teacher education, this paper proposes two things. First, training programs that focus entirely on policy enforcement risk producing instructors who can penalize cheating but are not prepared to restore learning. Second, one remedy is case-based professional training, in which teachers analyze real-world delegation incidents to practice identifying lost acquisitional opportunities and designing recovery strategies. Both are proposals derived from the framework, not findings. They are consistent with recent calls for responsible AI integration in language teacher development (Pham & Renandya, 2026), but their effectiveness remains to be established.

Limitations

This paper's argument has limits. The conceptual framework and claims about acquisition still need empirical testing, as proposed in Section 5. Until those studies are done, the argument remains a defensible prediction, not yet a proven result. The construct is also set out for the limiting case of full delegation. Its application to the intermediate positions on the continuum depends on the gradient in Proposition 2 holding, which is not yet established. The paper also leans on production-based SLA, mainly the output hypothesis and noticing accounts, which not everyone in the field of SLA accepts. A reader who rejects the output and noticing accounts may judge the cost of delegation differently. However, the main structural point still holds. Full delegation removes learner production.

Conclusion

Agentic displacement refers to a specific loss in L2 writing. It occurs when a learner delegates a writing task to an autonomous AI browser before doing the work of planning, drafting, monitoring, or revising. The self-regulated writing cycle never runs. Contract cheating already showed this loss in another form. What has changed is the route: a browser-based agent removes the delay, cost, and second human that once made outsourcing harder.

This paper defines agentic displacement, distinguishes it from cognitive offloading, contract cheating, and hybrid regulation, traces its consequences through production-based accounts of acquisition, states five testable propositions, and derives three dimensions of teacher knowledge. The framework is offered ahead of the evidence, deliberately, because the tools are already in learners' hands and the studies do not yet exist. The propositions in Section 5 specify how the framework can be wrong, and the boundary conditions in Section 3.3 specify where it may not apply. Both specifications are what entitle a conceptual construct to empirical attention.

For practice, the analysis relocates the response from detection to design and from policy to development. When an L2 learner hands a writing task to an autonomous agent, the learner does not plan, draft, notice their own interlanguage, or struggle with language under

communicative pressure. For a learner who is writing to acquire the language, those are not consequences for a grade. They are consequences for language.

Declarations

Declaration of generative AI use. During the preparation of this manuscript, generative AI tools, including Claude and ChatGPT, were used for literature organization, language refinement, structural checking, and editorial revision. The author retained responsibility for the conceptual argument, source selection, interpretation, verification of all references against their sources, and final wording.

Compliance with Ethical Standards

Funding. This research received no external funding.

Conflict of interest. The author declares no conflict of interest.

Research involving human participants and/or animals. This is a conceptual paper. No human participants or animals were involved, and no data were collected.

Informed consent. Not applicable.

Data availability. No datasets were generated or analyzed for this conceptual paper.

Author Contributions

The sole author was responsible for conceptualization, methodology, literature synthesis, writing of the original draft, review and editing, and final approval of the manuscript.

References

- Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher, R. W. Pew, L. M. Hough, & J. R. Pomerantz (Eds.), *Psychology and the real world: Essays illustrating fundamental contributions to society* (pp. 56–64). Worth Publishers.
- Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe & A. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185–205). MIT Press. <https://doi.org/10.7551/mitpress/4561.003.0011>
- Bozkurt, A., Crompton, H., & Fell Kurban, C. (2026). The devil is in the det[ai]ls: AI agents, ghost students, and the crisis of verified presence in an agentic AI world [Editorial]. *Open Praxis*, 18(1), 1–12. <https://doi.org/10.55982/openpraxis.18.1.1145>
- Bretag, T., Harper, R., Burton, M., Ellis, C., Newton, P., Rozenberg, P., Saddiqui, S., & van Haeringen, K. (2019). Contract cheating: A survey of Australian university students. *Studies in Higher Education*, 44(11), 1837–1856. <https://doi.org/10.1080/03075079.2018.1462788>
- California Virtual Campus. (2025, November). *Safeguarding academic integrity in the age of agentic AI* [White paper]. California Community Colleges. <https://www.cvc.edu/wp-content/uploads/2025/11/Safeguarding-Academic-Integrity-in-the-Age-of-Agentic-AI.pdf>

- Dai, L., Jiang, Y.-H., Chen, Y., Guo, Z., Liu, T.-Y., & Shao, X. (2025). Agent4EDU: Advancing AI for education with agentic workflows. In *Proceedings of the 2024 3rd International Conference on Artificial Intelligence and Education (ICAIE '24)* (pp. 180–185). Association for Computing Machinery.
<https://doi.org/10.1145/3722237.3722268>
- Dau, T. T. L., & Lam, T. T. (2026). Assessing L2 student writing in the AI era: A systematic review on challenges and best practices. *International Journal of TESOL & Education*, 6(1), 72–93. <https://doi.org/10.54855/ijte.26615>
- Dwivedi, Y. K., Helal, M. Y. I., Elgendy, I. A., Alahmad, R., Walton, P., Suh, A., Singh, V., & Jeon, I. (2026). Agentic AI systems: What it is and isn't. *Global Business and Organizational Excellence*, 45(3), 253–263. <https://doi.org/10.1002/joe.70018>
- Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2), 489–530. <https://doi.org/10.1111/bjet.13544>
- Feng, H., Li, K., & Zhang, L. J. (2025). What does AI bring to second language writing? A systematic review (2014–2024). *Language Learning & Technology*, 29(1), 1–27. <https://doi.org/10.64152/10125/73629>
- Garcia, P. H. S., de Souza, P. S. S., Luizelli, M. C., Mombach, J. G., Sachete, A. D. S., & Rossi, F. D. (2026). The new web attack surface: A taxonomy of semantic and agentic threats in AI browsers. *IEEE Access*, 14, 54109–54126. <https://doi.org/10.1109/ACCESS.2026.3680147>
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), Article 6. <https://doi.org/10.3390/soc15010006>
- Hou, C., Zhu, G., & Sudarshan, V. (2025). The role of critical thinking on undergraduates' reliance behaviours on generative AI in problem-solving. *British Journal of Educational Technology*, 56(5), 1919–1941. <https://doi.org/10.1111/bjet.13613>
- Hyland, K. (2026). Writing in the AI era: Rethinking writing, research and teaching. *Journal of Second Language Writing*, 72, Article 101302. <https://doi.org/10.1016/j.jslw.2026.101302>
- Jaakkola, E. (2020). Designing conceptual articles: Four approaches. *AMS Review*, 10(1–2), 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- Järvelä, S., Nguyen, A., & Hadwin, A. (2023). Human and artificial intelligence collaboration for socially shared regulation in learning. *British Journal of Educational Technology*, 54(5), 1057–1076. <https://doi.org/10.1111/bjet.13325>
- Jiang, F. K., & Hyland, K. (2025). Does ChatGPT write like a student? Engagement markers in argumentative essays. *Written Communication*. Advance online publication. <https://doi.org/10.1177/07410883251328311>
- Kofinas, A. K., Tsay, C. H.-H., & Pike, D. (2025). The impact of generative AI on academic integrity of authentic assessments within a higher education context. *British Journal of Educational Technology*, 56(6), 2522–2549. <https://doi.org/10.1111/bjet.13585>

- Kostopoulos, G., Gkamas, V., Rigou, M., & Kotsiantis, S. (2025). Agentic AI in education: State of the art and future directions. *IEEE Access*, 13, 177467–177491. <https://doi.org/10.1109/ACCESS.2025.3620473>
- Lee, S.-M. (2020). The impact of using machine translation on EFL students' writing. *Computer Assisted Language Learning*, 33(3), 157–175. <https://doi.org/10.1080/09588221.2018.1553186>
- Manchón, R. M. (Ed.). (2011). *Learning-to-write and writing-to-learn in an additional language*. John Benjamins. <https://doi.org/10.1075/llt.31>
- Microsoft. (n.d.). *Copilot in Edge*. Retrieved September 15, 2026, from <https://www.microsoft.com/en-us/edge/copilot-mode>
- Molenaar, I. (2022). The concept of hybrid human-AI regulation: Exemplifying how to support young learners' self-regulated learning. *Computers and Education: Artificial Intelligence*, 3, Article 100070. <https://doi.org/10.1016/j.caeai.2022.100070>
- Newton, P. M. (2018). How common is commercial contract cheating in higher education and is it increasing? A systematic review. *Frontiers in Education*, 3, Article 67. <https://doi.org/10.3389/feduc.2018.00067>
- OECD. (2026). *The agentic AI landscape and its conceptual foundations* (OECD Artificial Intelligence Papers, No. 56). OECD Publishing. <https://doi.org/10.1787/396cf758-en>
- OpenAI. (n.d.-a). *Evolving Atlas into ChatGPT for browser-based agentic work*. Retrieved September 15, 2026, from <https://help.openai.com/en/articles/20001371-evolving-atlas-into-chatgpt-for-browser-based-agentic-work>
- OpenAI. (n.d.-b). *Using Ask ChatGPT sidebar and ChatGPT agent on Atlas*. Retrieved September 12, 2026, from <https://help.openai.com/en/articles/12628199-using-ask-chatgpt-sidebar-and-chatgpt-agent-on-atlas>
- Perplexity. (2026, September 7). *Managing Comet Assistant permissions*. <https://www.perplexity.ai/help-center/en/articles/13531023-managing-comet-assistant-permissions>
- Pham, V. P. H., & Renandya, W. A. (2026). Navigating the AI turn: Framework for responsible AI integration in language education. *International Journal of TESOL & Education*, 6(2), 70–85. <https://doi.org/10.54855/ijte.26626>
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Roitman, A. (2025, October 22). *Cheating beyond ChatGPT: Agentic browsers present risks to universities* [Blog post]. Pangram Labs. <https://www.pangram.com/blog/agentic-browsers-present-risks-to-universities>
- Salah, M., Al-Busaidi, R., Al-Ghawi, A., & Shawky, W. (2026). The regulation paradox: Agentic AI, bounded autonomy, and self-regulated learning in higher education. *Frontiers in Education*, 11, Article 1902833. <https://doi.org/10.3389/feduc.2026.1902833>
- Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2026). AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges. *Information Fusion*, 126, Article 103599. <https://doi.org/10.1016/j.inffus.2025.103599>

- Schmidt, R. W. (1990). The role of consciousness in second language learning. *Applied Linguistics*, 11(2), 129–158. <https://doi.org/10.1093/applin/11.2.129>
- Swain, M. (1985). Communicative competence: Some roles of comprehensible input and comprehensible output in its development. In S. Gass & C. Madden (Eds.), *Input in second language acquisition* (pp. 235–253). Newbury House.
- Swain, M. (1995). Three functions of output in second language learning. In G. Cook & B. Seidlhofer (Eds.), *Principle and practice in applied linguistics: Studies in honour of H. G. Widdowson* (pp. 125–144). Oxford University Press.
- Swain, M., & Lapkin, S. (2002). Talking it through: Two French immersion learners' response to reformulation. *International Journal of Educational Research*, 37(3–4), 285–304. [https://doi.org/10.1016/S0883-0355\(03\)00006-5](https://doi.org/10.1016/S0883-0355(03)00006-5)
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4
- Teng, L. S. (2024). Individual differences in self-regulated learning: Exploring the nexus of motivational beliefs, self-efficacy, and SRL strategies in EFL writing. *Language Teaching Research*, 28(2), 366–388. <https://doi.org/10.1177/13621688211006881>
- Teng, L. S., & Zhang, L. J. (2016). A questionnaire-based validation of multidimensional models of self-regulated learning strategies. *The Modern Language Journal*, 100(3), 674–701. <https://doi.org/10.1111/modl.12339>
- Teng, L. S., & Zhang, L. J. (2020). Empowering learners in the second/foreign language classroom: Can self-regulated learning strategies-based writing instruction make a difference? *Journal of Second Language Writing*, 48, Article 100701. <https://doi.org/10.1016/j.jslw.2019.100701>
- Truscott, J. (1998). Noticing in second language acquisition: A critical review. *Second Language Research*, 14(2), 103–135. <https://doi.org/10.1191/026765898674803209>
- Vu, L. U., Tran, N. M., Le, T. K. H., & Dao, H. L. (2022). Applying writing feedback orientation and self-regulated learning writing strategies to EFL students at Van Lang University during COVID-19. *International Journal of TESOL & Education*, 2(5), 64–88. <https://doi.org/10.54855/ijte.22255>
- Warschauer, M., Tseng, W., Yim, S., Webster, T., Jacob, S., Du, Q., & Tate, T. (2023). The affordances and contradictions of AI-generated text for writers of English as a second or foreign language. *Journal of Second Language Writing*, 62, Article 101071. <https://doi.org/10.1016/j.jslw.2023.101071>
- Wiederhold, B. K. (2026). From human web to agentic web: When machines become the internet's primary users [Editorial]. *Cyberpsychology, Behavior, and Social Networking*. Advance online publication. <https://doi.org/10.1177/21522715261432850>
- Zhang, Y., Li, M., Chieh, Y., & Han, S. (2022). On a path to becoming more self-regulated: Reflective journals' impact on Chinese English as a foreign language students' self-regulated writing strategy use. *Frontiers in Psychology*, 13, Article 1042031. <https://doi.org/10.3389/fpsyg.2022.1042031>
- Zimmerman, B. J. (2000). Attaining self-regulation: A social cognitive perspective. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of self-regulation* (pp. 13–39). Academic Press. <https://doi.org/10.1016/B978-012109890-2/50031-7>

Zimmerman, B. J., & Risemberg, R. (1997). Becoming a self-regulated writer: A social cognitive perspective. *Contemporary Educational Psychology*, 22(1), 73–101.
<https://doi.org/10.1006/ceps.1997.0919>

Biodata

Hala Fawzi is an Assistant Professor in the English Language Department at Mohammed Jaber Al Ansari College for Teachers, University of Bahrain. Her research interests include AI in education, EFL/ESL writing, self-regulated learning, language teacher education, and responsible AI integration in English language teaching.