Aeterna Institute • De Litteris • I

De Suspensione

Inflection, Composition, and the Long Drift Outward

Rumi Allbert[1]

Aeterna Institute • May 27, 2026

Abstract. Three otherwise-heterogeneous historical movements share a common object: the Latin-to-Romance simplification, the digital externalization of memory, and the offloading of sentence production to generative artificial intelligence. That common object, here named *suspensio*, is the inner act of composition, the specific mental work of holding thought in suspension while form catches up with it. The practice was trained, in classical and medieval pedagogy, by the disciplines of the Trivium and by grammatical languages whose case endings obliged the speaker to track relation internally. Drawing on four literatures (classical-language impact studies, Romance morphosyntactic typology, neuro- and psycholinguistic work on grammar and thought, and the recent empirical work on cognitive offloading and large language models), this review defends a staged argument: inflection trains habits of attention that simplification redistributes rather than erases; digital retrieval externalizes content; and generative AI, for the first time, begins to externalize composition itself. The strong Whorfian thesis that linguistic simplification produced cultural shallowness is rejected on the evidence. The narrower and more pressing claim that survives is that three independent migrations of form, content, and production now converge on the same inner practice, and that recent proxy measures of compositional recall, ownership, effort, and offloading suggest the practice can thin under AI assistance.

Keywords: *suspensio*, *memoria*, the inner act of composition, cognitive offloading, large language models, morphosyntactic change, linguistic relativity, the Trivium, externalization, Aeterna Institute.

[1]Founder & Researcher, Aeterna Institute.

τοῦτο γὰρ τῶν μαθόντων λήθην μὲν ἐν ψυχαῖς παρέξει μνήμης ἀμελετησίᾳ, ἅτε διὰ πίστιν γραφῆς ἔξωθεν ὑπ᾽ ἀλλοτρίων τύπων, οὐκ ἔνδοθεν αὐτοὺς ὑφ᾽ αὑτῶν ἀναμιμνῃσκομένους.

This discovery of yours will create forgetfulness in the learners' souls, because they will not use their memories; they will trust to the external written characters and not remember of themselves.

— Plato, *Phaedrus* 275a (Plato, 1995)

♦ ♦ ♦

τοσοῦτον διαφέρουσιν οἱ πεπαιδευμένοι τῶν ἀπαιδεύτων ὅσον οἱ ζῶντες τῶν τεθνεώτων.

The educated differ from the uneducated as much as the living from the dead.

— Aristotle, reported in Diogenes Laertius V.19 (cf. *Protrepticus* fr.)

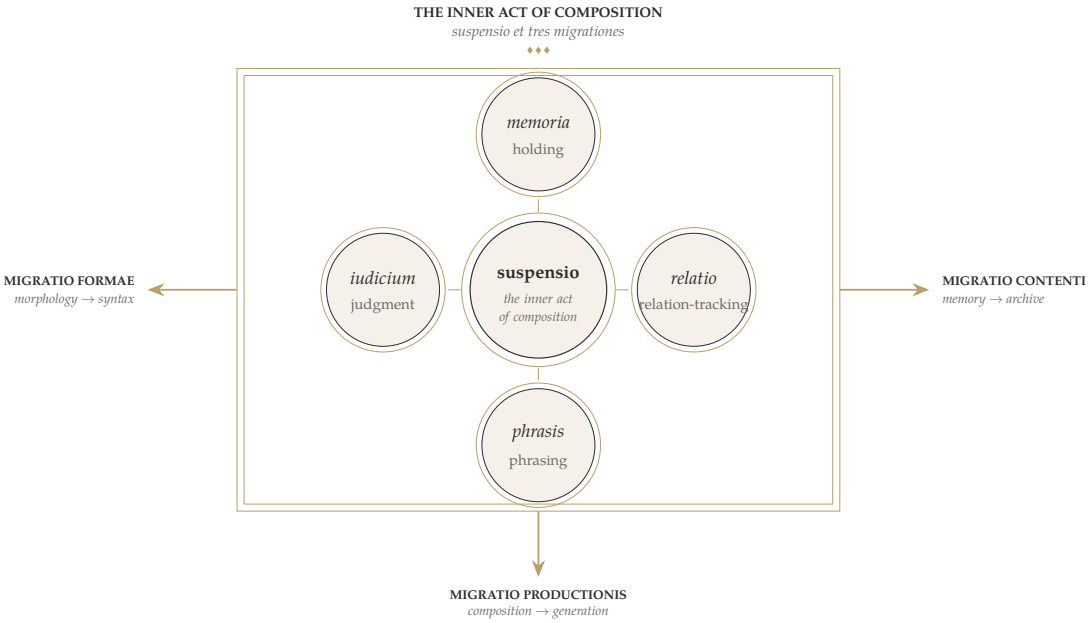

Figure 1. The paper's central claim in schematic form. At the centre lies the practice defended here: *suspensio*, the inner act of composition, the specific mental work of holding thought in suspension while form catches up with it. Around it are the four operations of the rational mind on which the practice draws: *memoria*, *relatio*, *phrasis*, *iudicium*. Outside the inner frame, three arrows break outward: the three historical migrations of form (Latin to analytic Romance), of content (memory to archive), and of production (author to generator). Each externalizes something different on a different timescale; all three converge on the same inner practice, and all three ask less of it.

1. Introduction

Consider a scene. A subject sits in the MIT Media Lab under an EEG cap, in front of a screen showing an essay prompt and an open ChatGPT window. She writes for twenty minutes, using the model to draft. The text on her screen at the end is fluent, adequate, recognizably an essay. Minutes after she finishes, the experimenter asks her to quote, out loud, a single sentence from what she has just written. She cannot. Eighty-three percent of her cohort cannot (Kosmyna et al., 2025). The sentences are present; the act of having composed them is not.

The missing thing is not a memory fact and not a vocabulary item. It is the residue of an act: the specific mental work of having composed the sentence oneself, inside, by holding a thought in suspension while form caught up with it. When that work is done by a tool, there is no residue, because there is no act; the output is on the screen, but nothing was performed within the speaker. Call this inner act of composition by the Latin term the classical pedagogical tradition already used for it: *suspensio*. It is the willingness, and the skill, to keep thought open until form arrives to discharge it.

The worry about externalization is not new. Plato in the *Phaedrus* feared that writing would hollow out the memory it pretended to serve (Plato, 1995); every subsequent major technology of expression, from the codex to the printing press to the internet, has occasioned a version of the same worry (Carr, 2010; Ong, 1982). What is new is

not the worry but what is being externalized. The *Phaedrus* anxiety was about *memoria*, the classical storehouse of the soul; the recent psychological literature on cognitive offloading (Risko and Gilbert, 2016; Skulmowski, 2023; Sparrow et al., 2011) has documented that particular anxiety empirically over two decades, with a mixed but real replication record. Since 2023, a second wave of evidence has extended the work to large language models (Gerlich, 2025; Kosmyna et al., 2025; Lee et al., 2025; Stadler et al., 2024) and reports, with striking consistency, that the tool is beginning to externalize something the old worries did not reach: not what is stored but what is performed, not the contents of thought but the act of composing a sentence in the first place.

Beneath these two contemporary literatures lies a third and much older one: the typology of morphosyntactic change, which traces the long-running drift from richly inflected classical Latin to its more analytic Romance descendants (Ledgeway, 2012; Liu et al., 2025; Nijs et al., 2025). On the usual account, this is a story about language. Read alongside the offloading literature it is also a story about the same inner practice. What a highly inflected language asks of its speaker is the relational tracking that *suspensio* makes possible; what an analytic language asks is less of that and more of position and context. The cognitive load does not vanish, but it migrates. The Latin case ending was an exterior form that required interior work; the English position is an exterior form that does more of that work on the speaker's behalf. That is the first migration. The externalized archive of digital retrieval is the second. The externalized composition of generative AI is the third.

Three otherwise-heterogeneous movements (one morphological, one informational, one generative) share a common object, *suspensio*, and a common direction, outward. They are not the same thing happening at different scales; they are different things that nevertheless ask less of the same inner practice. The strong Whorfian thesis that linguistic simplification produced cultural shallowness is rejected here on the evidence. What the evidence does support is narrower and more pressing: that the practice underlying considered speech and considered thought is under pressure, that recent empirical work supplies proxy evidence for its attenuation under AI assistance, and that the three migrations arrive at the same object from different directions. Table 1 summarizes the four literatures on which the argument rests.[2]

[2]This is a narrative review, not a systematic one. Studies were selected for their conceptual relevance and their standing in their respective literatures; where effects are contested or under-replicated, this is noted in the text.

Evidence strand	Strongest supported claim	Main limit	Principal references
Classical-language impact studies	Modest gains in grammar, vocabulary, and language-related achievement; no reliable effect on general intelligence	Methodologically uneven; mostly quasi-experimental	Vereeck et al. (2025); Bracke and Bradshaw (2020); Haag and Stern (2000); Hauspie et al. (2024); Masciantonio (1977)
Latin-to-Romance morphosyntactic change	Redistribution of grammatical load from morphology to syntax, driven by contact	Reorganization, not decline; social causation underdetermines cognitive consequences	Ledgeway (2012); Liu et al. (2025); Nijs et al. (2025); Trudgill (2011); McWhorter (2007); Bentz and Winter (2013); Dale and Lupyan (2012)
Grammar and thought	Grammatical categories can tilt perception, event framing, and attention	Relativistic effects are real but task- and context-dependent	Morrow (1986); Wolff and Holmes (2011); Thierry (2016); Boroditsky (2001); Samuel et al. (2019); Tal and Arnon (2022)
Cognitive offloading and generative AI	Offloading preserves convenience while redistributing memory, effort, and ownership; early AI studies suggest that the pattern extends into composition	Google-effect literature has mixed replication; AI literature is very young and its operationalizations are indirect	Risko and Gilbert (2016); Sparrow et al. (2011); Storm et al. (2017); Gong and Yang (2024); Skulmowski (2023); Stadler et al. (2024); Lee et al. (2025); Gerlich (2025); Kosmyna et al. (2025)

Table 1. The four evidence strands synthesized in this review. Each row states what the literature genuinely supports, where its methodological limits lie, and which references anchor the claim.

Two clarifications are owed before we proceed. *Suspensio* is an interpretive term, not a variable named by the studies cited here. No study below measures it directly. What the empirical literature measures are adjacent phenomena (grammatical role-tracking demands, memory for self-produced sentences, perceived ownership, cognitive offloading, critical engagement) from which the condition of the inner act of composition is inferred. The argument is also asymmetrical across the three migrations: the first is typological and reconstructive, the second concerns the externalization of content in the offloading literature, and the third has the closest empirical proxies because AI assistance lets researchers observe what happens when drafting is partly delegated. The claim is therefore one of conceptual convergence, not of uniform measurement.

2. The Discipline of Suspension

De Suspensione

Translating a Latin sentence is not, in the first instance, a vocabulary problem. It is a *suspensio* problem. A noun appears before its role is fully determined. A verb arrives late, at the end of the clause, announcing the shape of everything that preceded it. Case endings and agreement oblige the reader to hold several possible parses alive, adjacent and undecided, until the morphology accumulates enough evidence to let resolution occur. The mental act this requires is neither parsing nor decoding. It is the voluntary deferral of closure: the willingness to keep thought open, suspended, in a condition of relational expectancy, until the form of the sentence arrives to discharge it. We call this practice *suspensio*: the inner act of composition in its receptive mode. It is one thing to read a sentence in which the work of relation has already been done by word order; it is another to perform that relational work oneself, inside, while the sentence unfolds.

The classical tradition already knew this practice, and trained it explicitly. Quintilian in the *Institutio Oratoria* I.1 lays out a curriculum in which the child's mind is drilled, morpheme by morpheme, on the recognition and deployment of grammatical relation; the discipline is not rote but attention. Cicero in *De Oratore* II distinguishes *meditatio* (the inward composition that precedes speech) from *memoria* (the cultivated storehouse that makes *meditatio* possible), a distinction that loses its meaning the moment composition is outsourced. Augustine, two centuries later, gives the phenomenology its deepest form: in *Confessions* XI he calls this holding-open the *distentio animi*, the stretching of the soul across the interval between what is remembered, what is present, and what is expected. All three are describing the same act under different aspects: the disciplined interiority that sentence-building, in a highly inflected language, makes unavoidable. Sister Miriam Joseph, in her twentieth-century recovery of the Trivium, names the three operations of the rational mind as simple apprehension, judgment, and discursive reasoning (Sister Miriam Joseph, 2002). *Suspensio* is not one of these operations but the connective tissue between them, the willingness to keep the first open until the second and third can be performed.

The modern empirical literature, approached from outside this tradition, arrives at the same object with weaker vocabulary. Vereeck et al. (2025) map the classical-language impact studies and conclude that the claim "Latin trains the mind" is real enough to be researchable but not yet well settled; Adema (2019) finds the field scarce and often driven more by teacher intuition than by systematic evidence. The better-controlled studies find modest and specific effects, not generalized cognitive enhancement. Haag and Stern (2000) report no meaningful difference in general intelligence between pupils who begin with Latin and those who begin with English; small advantages appear only on grammar-related tasks. Ortner et al. (2004) find no relevant effect on general reasoning after two years of Latin instruction. Bracke and Bradshaw (2020), reviewing the wider literature, reach a similar middle position: positive findings exist but cluster around language-related outcomes rather than general cogni-

tion. Hauspie et al. (2024) report that prior Latin study predicts later achievement in higher education, particularly in linguistic subjects, though the effect is neither broad nor uniform. Masciantonio (1977) surveys older educational projects claiming gains in English verbal skill among Latin students, but the methodology of those reports is by modern standards uneven.

The narrower picture is the one the argument needs. Classical languages do not, on the evidence, produce superior minds in general. What they do, according to the studies that survive replication, is cultivate specific forms of verbal analysis, delayed interpretation, and attention to structure: *suspensio* under its several names. Maharaj (2020) treats translation as a practice of formation rather than as a technical skill, an interpretive overlay on the empirical finding but one that captures the phenomenology well. *Suspensio* is not a curiosity of Latin pedagogy. It is the common object of the three migrations traced in the sections that follow, the single inner practice that each successive externalization asks the speaker to perform less of.

3. The First Migration

Migratio Formae

What left the word had to reappear somewhere. That is the simplest summary of the first migration, which carries *suspensio* out of the inflected sentence and into the surrounding apparatus of fixed order, prepositions, and auxiliary verbs. The historical literature is firm on the facts: Latin had richer nominal case marking and freer word order than any of its Romance daughters, and over time much of that morphology was reduced or lost while fixed constituent order and functional structure took over more of the grammatical work (Ledgeway, 2010, 2012). The literature is equally firm that this should not be framed as simple degeneration. What is less often said is that the typological trade-off has a cognitive correlate: the interior work that a Latin case ending demanded of its reader is work that an English positional constraint performs on the reader's behalf, before the interior has to hold it. *Suspensio* has migrated outward, into form.

Ledgeway (2010) reads the shift from Latin to Romance as a move from a more synthetic system toward more analytic structures: fewer endings doing grammatical work, more reliance on separate words and on tighter order. Liu et al. (2025) quantify the same long pattern across Romance and find support for a trade-off between morphological richness and word-order freedom. Nijs et al. (2025), examining five Western European languages, find that declines in morphological complexity tend to *precede* increases in word-order rigidity: when endings stop carrying enough information, syntax has to become more disciplined. Figure 3 displays this schematic relationship, and Figure 2 makes the qualitative difference between the two regimes (morphological versus positional) concrete in a single minimal pair.

The social explanation for this drift is well established. Trudgill (2011) argues that morphological complexity is sustained by small, tight-knit speech communities (what

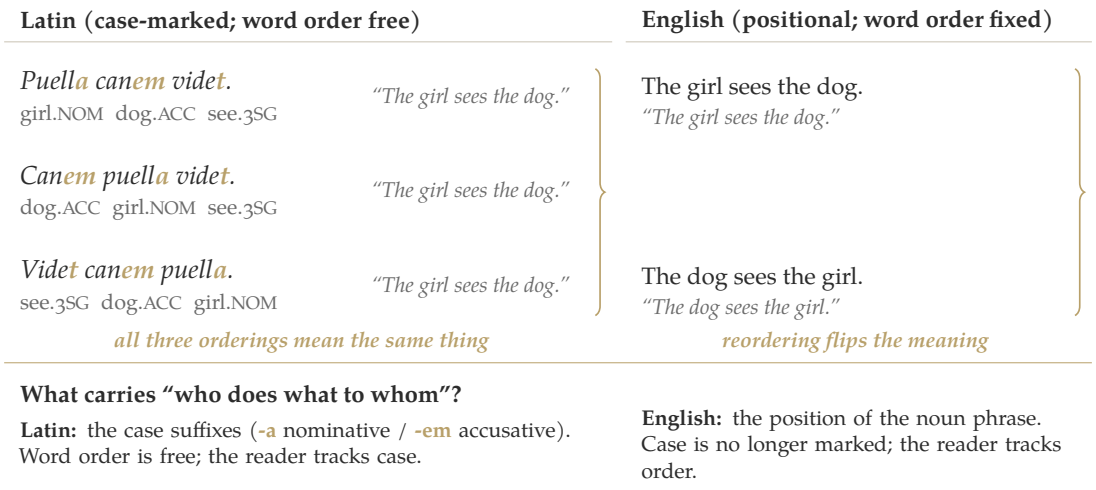

Figure 2. A minimal pair for the redistribution argument. In Latin, case suffixes (highlighted) carry grammatical role and the sentence may be reordered freely without loss of meaning; the cognitive work of role-tracking lives inside the morphology. In English, the same words rearranged produce a different sentence entirely; the work has migrated into position. This illustrates what §3 argues from the typological literature: what left the word had to reappear elsewhere.

he calls "societies of intimates") and simplifies when communities grow, disperse, and admit adult learners at scale. McWhorter (2007) provides the most-cited statement of the specific mechanism: sustained adult-L2 contact leaves measurable "analytic drift" even in languages that never creolized. Bentz and Winter (2013) show cross-linguistically that languages with more L2 learners tend to lose nominal case. Dale and Lupyan (2012) frame the same pattern under a "linguistic niche" hypothesis: large populations, adult learners, and frequent cross-group contact favour simplification of inflectional paradigms and greater transparency.

On this view, simplification is partly a democratic success: language becomes easier to learn across wider networks. The cost is not loss of meaning. The cost, more narrowly, is that less meaning is packed into the word itself. Rich morphology places more structure inside the utterance; analytic language spreads that structure across position, helper words, and context. Neither freely, the mind still interacts with structure, but with a different structure, one whose demands track different attentional habits.

4. Grammar as a Training of Attention

De Habitu Attentionis

Inflection matters, on the evidence, not because it is old, difficult, or noble, but because it externalizes less and demands more active internal tracking of relation. That is the moderate linguistic-relativity position, and it is the one this paper defends. The strong Whorfian form of the claim, that one language family makes its speakers globally smarter than another, is not supported by any current evidence, and was never the argument worth defending. Havelock (1963) reminded the twentieth century that literate rationality is itself a specific cognitive achievement rather than the default state

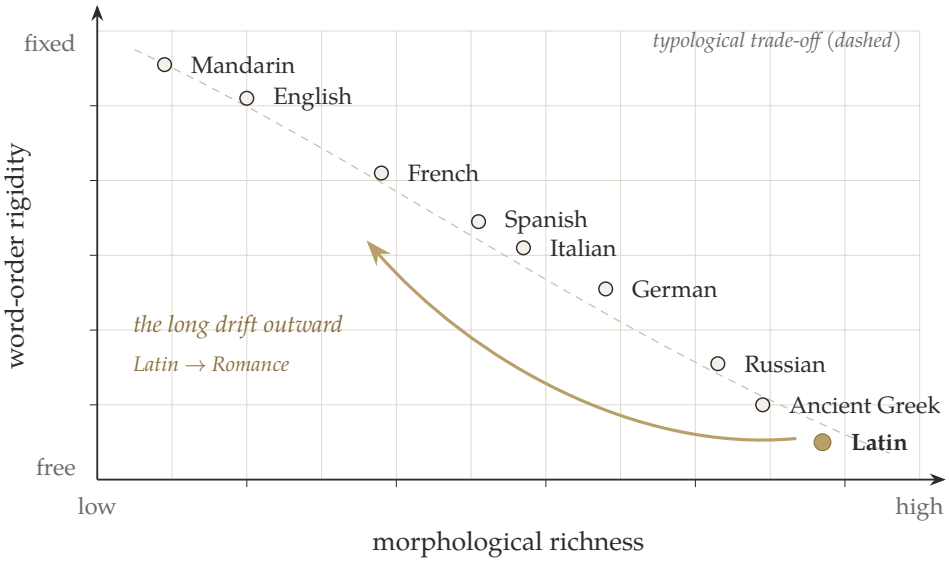

Figure 3. The morphology-and-syntax trade-off, with the Latin-to-Romance drift drawn onto it. Each circle places a language on two axes: horizontal, how much grammatical work is done by morphology (richness); vertical, how much is done by fixed word order (rigidity). Positions reproduce the qualitative ranking reported by Liu et al. (2025) for Romance and Nijs et al. (2025) for Western European languages more broadly; coordinates are indicative, not measurements. The gold arrow marks the historical drift named in §3: what left the inside of the word reappeared in the order between words.

of the mind, and the same caution now applies to grammatical form: what a given grammar makes easy to think, or hard to ignore, is a contingent training, not an imprisonment.

Morrow (1986) showed that grammatical morphemes help organize discourse processing by mapping conceptual distinctions into the flow of comprehension. Wolff and Holmes (2011) review the modern linguistic-relativity literature and explicitly reject strong determinism, but find evidence that language can make certain distinctions hard to ignore and can augment some forms of thinking. Thierry (2016) goes further in a neurolinguistic review, reporting that lexical and grammatical patterns can shape perception, event framing, semantic association, and aspects of executive function. Boroditsky (2001) provides a widely cited demonstration outside the European tradition, showing that Mandarin and English speakers, whose languages conventionally spatialize time differently, differ reliably in some temporal-reasoning tasks. That particular finding has had a contested replication record, and is cited here as illustrative of a pattern rather than as itself decisive.

The evidence is strongest where it is narrowest. Samuel et al. (2019), in a systematic review of forty-three empirical studies on grammatical gender, find that relativistic effects are real but task- and context-dependent, and that strong claims about gender-shaped cognition often have equally viable alternative explanations. The same caution applies here: grammar does not imprison the mind, but it does tilt perception and conceptualization in measurable ways.

Inflection matters, in this moderate form, because it externalizes less and demands more active internal tracking of relation. Case marking is a built-in guide to agency and dependence. Aspect trains attention to the contour of action. Mood and voice keep distinctions alive that analytic prose lets fade into paraphrase. Tal and Arnon (2022) show that redundant grammatical marking can aid learning: more marking is not always clutter; sometimes it gives the learner more handles on structure. Grammar is not just a delivery vehicle for thought; it is one of the recurring ways thought is disciplined. What cannot be said, on present evidence, is that the move from inflected to analytic systems straightforwardly caused cognitive decline.

5. The Second and Third Migrations

Migratio Contenti et Productionis

The sharpest contemporary evidence on externalization is not in the history of Latin; it is in the recent psychological literature on cognitive offloading and generative artificial intelligence. Here the two remaining migrations arrive: first the externalization of content, then the externalization of production, and with them the closest empirical proxies for *suspensio*. The inner act itself remains an interpretive construct, but some of its likely residues and conditions (recall of self-produced sentences, felt ownership, cognitive offloading, and the shift from composing to adjudicating) can be studied in laboratories, surveys, and mixed-methods designs.

A clarification is owed up front, because the AI literature attracts careless versions of the argument from every direction. The claim is not that using AI makes people dumber, nor that delegation to tools is itself a failure of mind; tools have always extended the mind, and nothing in what follows denies that. The claim is narrower. In its stronger uses, AI assistance now performs an act that used to be a constitutive part of being literate: the holding of a draft open against one's own judgment, while form catches up with thought. That act does not survive long when nobody is performing it. Faculties not exercised attenuate; what is at stake is not intelligence but a practice.

Risko and Gilbert (2016) define cognitive offloading as the use of the environment to reduce internal cognitive demand. The phenomenon is real: when people expect information to be available externally, they are less likely to remember the information itself and more likely to remember where to find it (Sparrow et al., 2011). The effect has had a mixed replication record (Storm et al., 2017), and a 2024 meta-analysis reports that "Google effects" on memory are real but moderated by perceived reliability of the external store, device type, and cognitive load (Gong and Yang, 2024). The cumulative picture is not one of a single dramatic effect; it is one of a distributed, domain-dependent shift in what is retained and what is offloaded. Skulmowski (2023) synthesizes this literature as a model of digital externalization: offloading does not leave the mind empty, but it leaves the mind holding biological pointers and gist traces in place of detail, and the paper warns of an *illusio cognitionis*, an illusion of knowing in which easy access is mistaken for mastery. This is the second migration. *Suspensio* is touched here only at the edges: the inner act of composition does not, in itself,

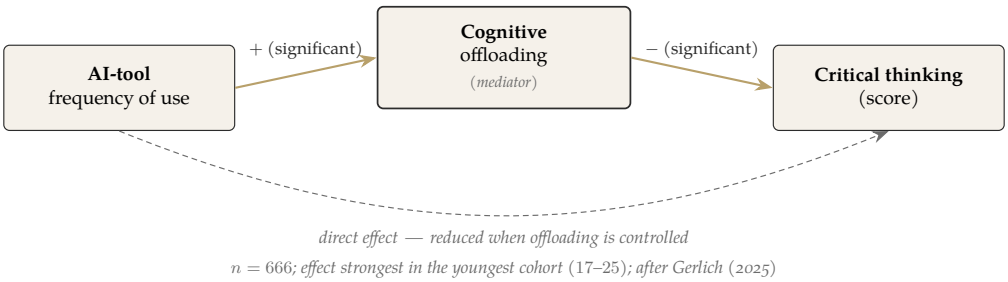

Figure 4. Mediation path diagram for the Gerlich (2025) finding. Frequency of AI-tool use predicts higher cognitive offloading, which in turn predicts lower critical-thinking scores. When the mediating variable is controlled, the direct path from AI use to critical thinking weakens substantially. Coefficient signs shown; the full standardised effect sizes are reported in the original paper.

require remembering every quoted sentence, only performing the act of composing around it. What thins, in this regime, is the ecology in which *suspensio* takes place: the half-remembered phrase that rises to meet the new sentence, the resonance, the allusion.

The third migration is qualitatively different. Generative AI does not merely store; in its stronger uses it proposes or performs substantial stretches of composition. Stadler et al. (2024) report that large language models reduce mental effort in student scientific inquiry while compromising depth. Lee et al. (2025), surveying knowledge workers, find that AI changes critical thinking rather than eliminating it: the worker is shifted from composing toward adjudicating. Confidence in the AI is the strongest moderator of whether engagement rises or falls. That is to say: the very *suspensio*-shaped act of holding one's own draft open against one's own judgment is liable to be replaced by the more frictionless acceptance or rejection of an externally generated one. The empirical pattern here recalls Pieper (1998), whose distinction between *ratio* (the discursive labor of composition and judgment) and *intellectus* (the receptive contemplation that discursive labor serves) maps directly onto the concern: AI-assisted work threatens to collapse *ratio* into the mere adjudication of generated outputs, leaving neither the labor nor what the labor is for. The most quantitatively robust evidence currently available is Gerlich (2025), a mixed-methods study of 666 participants that finds a significant negative correlation between frequency of AI-tool use and critical-thinking scores, mediated by increased cognitive offloading, with the effect strongest among younger users. Figure 4 formalizes that mediation as a path diagram.

Kosmyna et al. (2025) provide the most vivid, though still preliminary, evidence. In an essay-writing task with EEG monitoring and n=54 at baseline (n=18 at follow-up), participants using ChatGPT showed weaker brain connectivity than brain-only writers, lower ownership of their texts, and striking difficulty quoting sentences they had just produced; 83% of LLM-assisted writers could not reproduce a sentence they had written minutes earlier. The paper is a preprint and its own authors resist generalization, but its image is consistent with what the larger Gerlich data and the Lee survey already indicate: when the tool helps too much, authorship can thin, and the text may

exist with a weaker interior trace of having been composed.

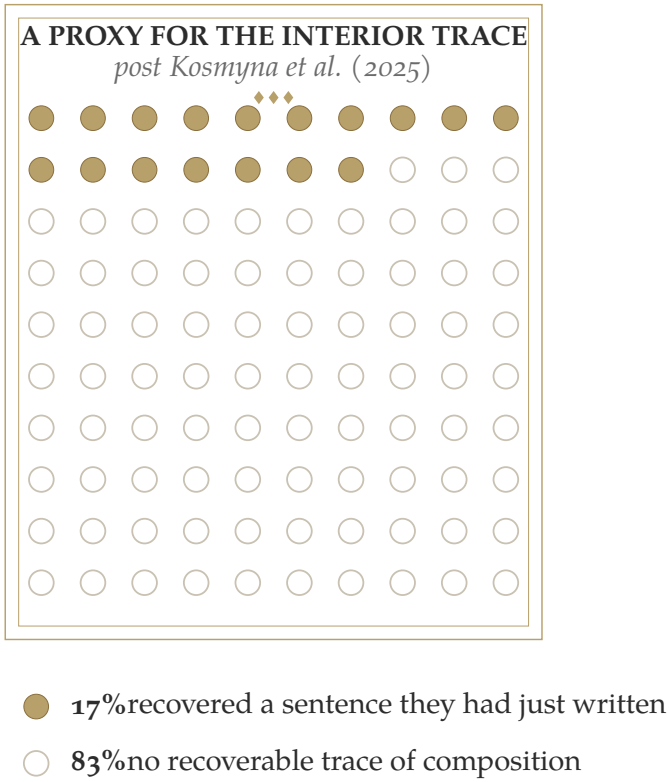

Figure 5. A proxy measure, not a direct measure, of *suspensio*. In Kosmyna et al. (2025), an EEG-monitored essay-writing task with $n = 54$ at baseline, 83% of LLM-assisted writers could not reproduce a sentence they had composed minutes earlier; the sentences were present on their screens, but many participants could not immediately recover a line as their own. Each dot is one percent; the figure renders the reported proportion, not the raw count. The inference drawn in this paper is interpretive: failure to reproduce a sentence one has just produced is treated here as one sign of a weakened interior trace, not as a direct operationalization of *suspensio* itself.

The pattern across these four studies is worth noting in itself, because no one of them could carry the claim alone. Gerlich (2025) is a large-N peer-reviewed mixed-methods survey of 666 participants: correlational, generalizable, but unable on its own to ground a causal inference. Kosmyna et al. (2025) is a small EEG-monitored experimental preprint of 54: invasive, vivid, and at the wrong end of every generalizability dimension. Lee et al. (2025) is a CHI-published self-report survey of knowledge workers: peer-reviewed and ecologically valid, but vulnerable to the usual self-report distortions. Stadler et al. (2024) is a controlled experimental paradigm in student scientific inquiry: peer-reviewed and causally clean, but narrow in domain. Each study's weakness is another study's strength. What gives the claim its evidential weight is not any single effect size but the fact that four independent literatures, with four different methodologies, converge on the same directional finding: under AI assistance, the internal trace of composition is attenuated.

Figure 6 is a heuristic rather than an estimate. It arranges the regimes on which the

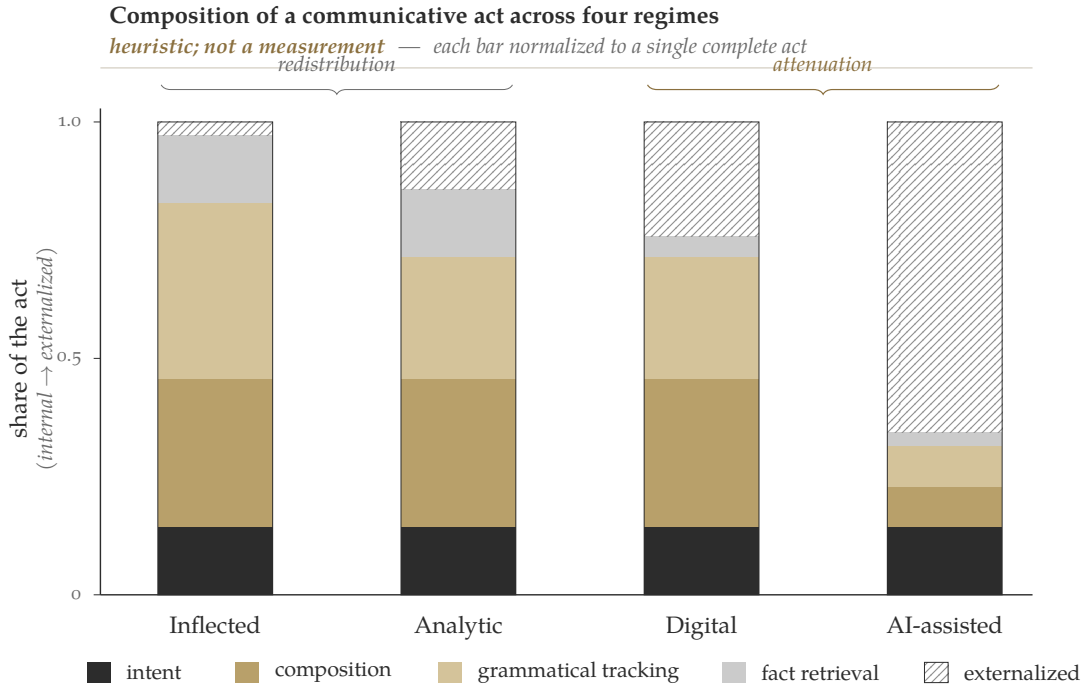

Figure 6. Each bar represents a complete communicative act, normalized to the same total height; the visual claim is therefore about *composition*, not absolute cognitive magnitude. The shift from Inflected to Analytic is a *redistribution*: grammatical work moves out of morphology and into syntax and context, but the share of the act performed inside the speaker stays roughly constant. The shift from Digital to AI-assisted is an *attenuation*: as composition itself is partly delegated, the hatched "externalized" region grows and the bottom three internal layers thin. The figure is a heuristic summary of claims argued in the surrounding text; it is not, and is not meant to be, a measurement. Sources for the underlying directional claims are Gerlich (2025); Gong and Yang (2024); Kosmyna et al. (2025); Lee et al. (2025); Risko and Gilbert (2016); Sparrow et al. (2011); Stadler et al. (2024).

argument turns as stacked columns and asks what portion of each communicative act remains internal to the speaker. The values are illustrative, not measurements, and the figure should be read as a synopsis of claims argued elsewhere in the text rather than as an independent result.

6. The Three Migrations

Tres Migrationes

Across the Latin-to-Romance simplification, the digital externalization of memory, and the offloading of sentence production to generative AI, the same inner practice (*suspensio*, the inner act of composition) is asked to bear less of the communicative act. Together the three movements are the *Tres Migrationes*: the migrations of form, of content, and of production.

They are not points on a single continuum, nor are they equally documented. The first is a redistribution of grammatical work; the second alters the mnemonic ecology in which composition takes place; the third can move part of composition itself outside the speaker. Their social meanings differ and their evidence bases differ. What they

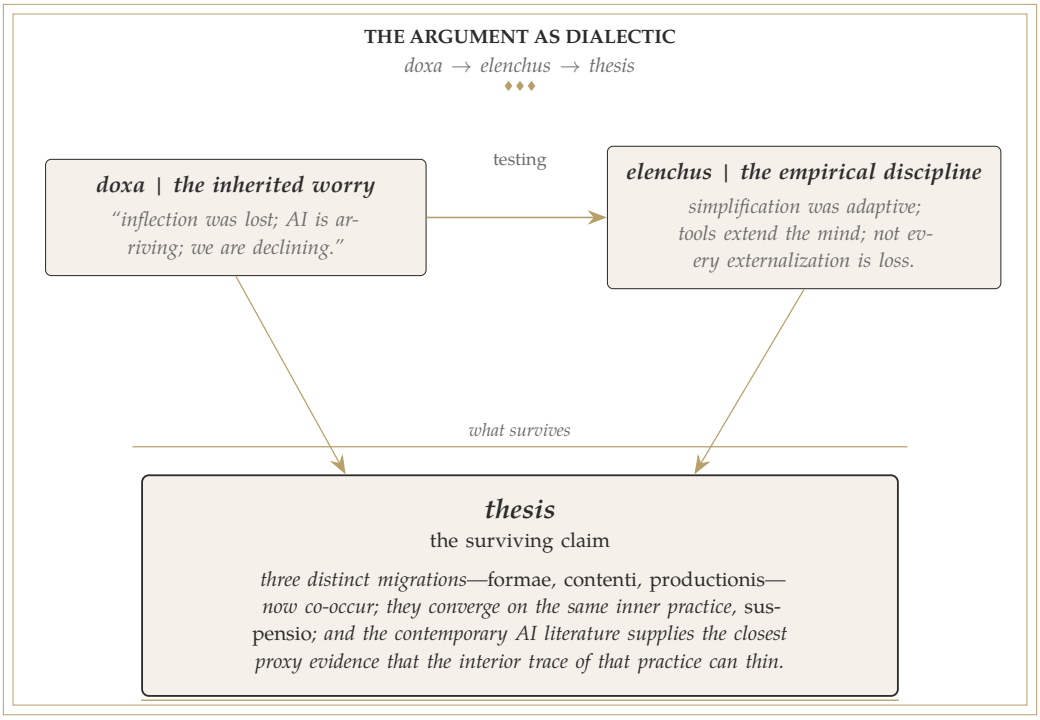

Figure 7. The paper's argument in the shape of its method. What the popular story reports (*doxa*) and what the empirical literature refuses to say (*elenchus*) are both true, and neither is the whole claim. The position this paper defends is the thesis that survives the encounter of the two: three distinct migrations of form, content, and production now co-occur, they converge on the same inner practice (*suspensio*), and the contemporary AI literature offers the closest proxy evidence that the interior trace of that practice can thin. The figure is not itself evidence; it makes explicit the argumentative discipline by which the paper narrows a cultural worry into a defensible claim.

share is a family resemblance: more infrastructure outside the speaker, less obligatory work inside. That family resemblance is the central claim of this review. Figure 7 sets out the argumentative shape by which the claim is reached.

The honest form of the cultural worry, then, is not that Latin was lost, nor that AI is arriving, nor that any single migration is a catastrophe. It is that *all three migrations are now co-occurring*, and that the inner practice which remains after all three have moved outward is no longer obvious, no longer performed by default, and no longer visible to the person from whom it has migrated. That is a cultural claim, not a linguistic or a neurological one, and it is defended here on cultural grounds.

Figure 8 places the intellectual history of the worry on a single timeline, a genealogy of recurring anxieties about externalization rather than evidence for the present claim. What changes between one voicing and the next is which dimension of the cognitive act is being moved outside.

7. Objections

Four objections deserve explicit reply.

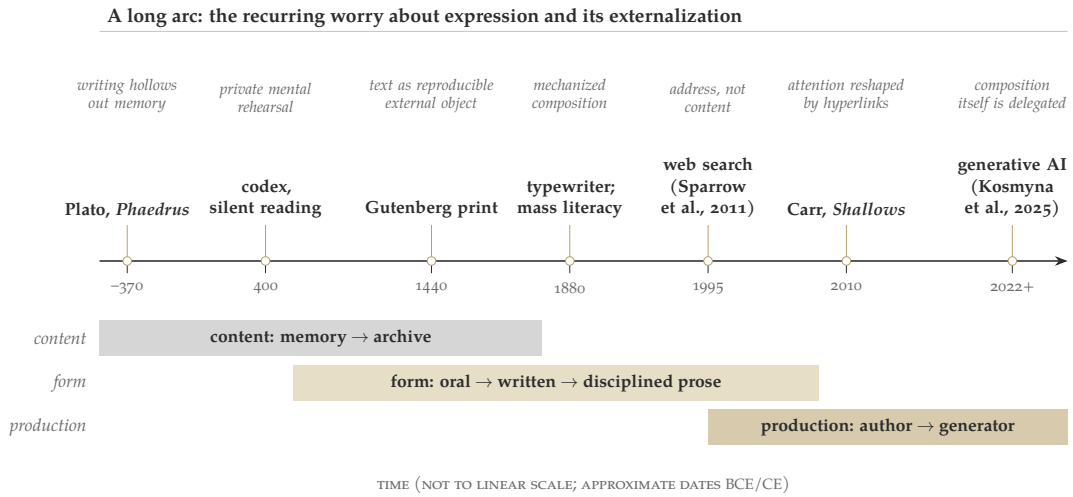

Figure 8. A compressed timeline of the recurrent worry that new technologies of expression externalize some part of the cognitive act. Upper nodes are representative thinkers and events; lower horizontal strips indicate the dimension of externalization that enters the argument at each stage. The timeline is not linear: the intent is to show *which dimension* was salient when, not to scale historical durations or to serve as independent proof of the paper's thesis.

Simpler languages are a democratic gain, not a loss. Partly true, and the paper grants it. Simplification of inflectional paradigms is closely tied to large speech communities, adult learners, and cross-group contact (Bentz and Winter, 2013; Dale and Lupyan, 2012; McWhorter, 2007; Trudgill, 2011); analytic structure makes a language easier to acquire past childhood. The paper's claim is not that the shift was bad (it often was not), but that it was a redistribution. What left the word reappeared elsewhere. The democratic gain is real; so is the migration of *suspensio* outward from morphology into position and context. Figure 3 is precisely that redistribution, not a decline.

Tools have always extended the mind. True, and the argument is not against extension as such. What is distinctive about the present moment is specific: the extended system now drafts the sentences, not only stores or retrieves them. Three independent methodologies converge on the empirical point that this particular extension changes the internal trace of the act: Kosmyna's EEG, Stadler's experimental paradigm, and Gerlich's mixed-methods survey (Gerlich, 2025; Kosmyna et al., 2025; Stadler et al., 2024). The argument turns on which part of the mind this particular tool extends past, and on what remains to be performed when it has.

The offloading literature is less settled than the paper implies. A fair objection. The original "Google effect" (Sparrow et al., 2011) has a mixed replication record (Storm et al., 2017); the 2024 meta-analysis reports the phenomenon as real but strongly moderated (Gong and Yang, 2024). The AI-effect literature is young: Kosmyna et al. (2025) is a preprint; Lee et al. (2025) is self-report; Stadler et al. (2024) is a single experimental paradigm. The strongest single data point is now Gerlich (2025), and although it is peer-reviewed and large-n, a correlational survey cannot by itself ground causal claims. The paper's argument does not require any single study; it rests on the directional agreement of four independent literatures. The empirical picture is young; the

long arc is not.

The thesis risks nostalgia. The objection deserves an answer but not the defensive one. The paper does not claim that Latin was good and English is bad; it claims that *suspensio* is a practice, and that practices, in the sense MacIntyre (1981) gives the word, are traditions of making that require communities to sustain. When no one performs a practice, the practice does not merely become harder to find. It ceases to exist as a practice, because the practice is the performing. What the paper defends is not classical pedagogy as such but the conditions under which the inner act of composition can still be trained. That those conditions have historically been classical ones is a contingent and, under the pressures of the three migrations, an increasingly important fact.

8. Conclusion

These four literatures land on a single result. Classical-language study defends itself, on the evidence, as a training of specific habits of verbal analysis and attention to structure, though not of generalized cognitive capacity (Bracke and Bradshaw, 2020; Haag and Stern, 2000; Hauspie et al., 2024; Vereeck et al., 2025). The Latin-to-Romance simplification is a redistribution of grammatical load from the inside of the word to the outside of the sentence, not a decay (Ledgeway, 2012; Liu et al., 2025; McWhorter, 2007; Nijs et al., 2025; Trudgill, 2011). The moderate linguistic-relativity position is that grammatical systems train what a speaker must notice, hold, and discriminate, without imprisoning thought (Morrow, 1986; Samuel et al., 2019; Thierry, 2016; Wolff and Holmes, 2011). And the contemporary evidence on AI-assisted composition shows, consistently across four independent methodologies, that under conditions of tool use we now have proxy evidence for an attenuation of the interior trace of composition (Gerlich, 2025; Kosmyna et al., 2025; Lee et al., 2025; Skulmowski, 2023). Read together, these four literatures converge on the same object. Each of the three migrations this paper has traced (of form, of content, of production) asks less of the same inner practice, and that practice, the one we have called *suspensio*, is the common object.

Suspensio leaves no artifact. It is legible only in its exercise; it produces no object one could point to outside the act of composing, holding, or deferring. This is why the Kosmyna finding is so unsettling. What was measured there was not *suspensio* itself, but one possible residue of it: whether a sentence produced minutes earlier had left enough of an interior trace to be recovered as one's own. For eighty-three percent of the LLM-assisted group, it had not (Kosmyna et al., 2025). The practice in question does not fail visibly. It simply leaves a weaker residue, and its attenuation is inferred from that weakness rather than directly observed.

Convenience is not neutral. Every gain in ease raises the question of what inner faculty is no longer being exercised, and that question does, at least partially, answer itself. Faculties not exercised attenuate. The faculty in question here is the one on which

considered speech, considered judgment, and considered thought all depend: not case endings, not memory palaces, not even the sentence as such, but the willingness to hold thought open until form arrives to discharge it. A civilization that externalizes this practice across all three dimensions at once has not lost Latin. It risks losing the condition under which Latin was worth having.

References

Suzanne Adema. Latin learning and instruction as a research field. *Journal of Latin Linguistics*, 18:35–59, 2019. doi: 10.1515/joll-2019-0001.

Christian Bentz and Bodo Winter. Languages with more second language learners tend to lose nominal case. *Language Dynamics and Change*, 3:1–27, 2013. doi: 10.1163/22105832-13030105.

Lera Boroditsky. Does language shape thought?: Mandarin and English speakers' conceptions of time. *Cognitive Psychology*, 43:1–22, 2001. doi: 10.1006/cogp.2001.0748.

Evelien Bracke and Carrie Bradshaw. The impact of learning Latin on school pupils: A review of existing data. *The Language Learning Journal*, 48:226–236, 2020. doi: 10.1080/09571736.2017.1400578.

Nicholas Carr. *The Shallows: What the Internet Is Doing to Our Brains*. W. W. Norton, New York, 2010.

Rick Dale and Gary Lupyan. Understanding the origins of morphological diversity: The linguistic niche hypothesis. *Advances in Complex Systems*, 15, 2012. doi: 10.1142/S0219525911500172.

Michael Gerlich. AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1):6, 2025. doi: 10.3390/soc15010006.

Chenghui Gong and Ying Yang. Google effects on memory: A meta-analytical review of the media effects of intensive Internet search behavior. *Frontiers in Public Health*, 12:1332030, 2024. doi: 10.3389/fpubh.2024.1332030.

Ludwig Haag and Elsbeth Stern. Non scholae sed vitae discimus? Auf der Suche nach globalen und spezifischen Transfereffekten des Lateinunterrichts. *Zeitschrift für Pädagogische Psychologie*, 14:146–157, 2000. doi: 10.1024//1010-0652.14.23.146.

Cedric Hauspie, Stijn Schelfhout, Nicolas Dirix, Lot Fonteyne, Mark Janse, Arnaud Szmalec, Astrid Vereeck, and Wouter Duyck. Does studying Latin in secondary education predict study achievement in academic higher education? *Language Learning*, 74, 2024. doi: 10.1111/lang.12639.

Eric A. Havelock. *Preface to Plato*. Belknap Press of Harvard University Press, Cambridge, Mass., 1963.

Natalia Kosmyna, Elias Hauptmann, Yue Yuan, Jessica Situ, Xiuqi Liao, Alicia V. Beresnitzky, Iris Braunstein, and Pattie Maes. Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. *arXiv preprint arXiv:2506.08872*, 2025.

Adam Ledgeway. Syntactic and morphosyntactic typology and change. In *The Cambridge History of the Romance Languages*, pages 382–471. Cambridge University Press, 2010. doi: 10.1017/CHOL9780521800723.010.

Adam Ledgeway. *From Latin to Romance: Morphosyntactic Typology and Change*. Oxford University Press, 2012. doi: 10.1093/acprof:oso/9780199584376.001.0001.

Hao-Ping Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025. doi: 10.1145/3706598.3713778.

Song Liu, Jianwei Yan, and Haitao Liu. The complexity trade-off between morphological richness and word order freedom in Romance languages: A quantitative perspective. *Zeitschrift für Romanische Philologie*, 141:323–349, 2025. doi: 10.1515/zrp-2025-0032.

Alasdair MacIntyre. *After Virtue: A Study in Moral Theory*. University of Notre Dame Press, Notre Dame, IN, third edition, 1981.

Neil Maharaj. A phenomenological reflection on translating Latin: What education does with us. *Educational Studies*, 56:584–598, 2020. doi: 10.1080/00131946.2020.1837831.

Rudolph Masciantonio. Tangible benefits of the study of Latin: A review of research. *Foreign Language Annals*, 10:375–382, 1977. doi: 10.1111/j.1944-9720.1977.tb02999.x.

John H. McWhorter. *Language Interrupted: Signs of Non-Native Acquisition in Standard Language Grammars*. Oxford University Press, Oxford, 2007.

Daniel G. Morrow. Grammatical morphemes and conceptual structure in discourse processing. *Cognitive Science*, 10:423–455, 1986. doi: 10.1207/s15516709cog1004_2.

Jorik Nijs, Freek Van de Velde, and Hubert Cuyckens. Is word order responsive to morphology? Disentangling cause and effect in morphosyntactic change in five Western European languages. *Entropy*, 27, 2025. doi: 10.3390/e27010053.

Walter J. Ong. *Orality and Literacy: The Technologizing of the Word*. Methuen, London, 1982.

Tuulia M. Ortner, Miriam Asanger, Klaus D. Kubinger, and René T. Proyer. Zur Frage nach Auswirkungen von Lateinunterricht auf die kognitive Fähigkeit "Reasoning". *Psychologie in Erziehung und Unterricht*, 55:189–195, 2004. doi: 10.5167/UZH-3636.

Josef Pieper. *Leisure, the Basis of Culture*. St. Augustine's Press, South Bend, IN, 1998. First German edition: *Musse und Kult*, 1948; first English translation 1952; St. Augustine's edition used here.

Plato. *Phaedrus*. Hackett, Indianapolis, 1995. Stephanus pagination 274c–275b; orig. c. 370 BCE.

Evan F. Risko and Sam J. Gilbert. Cognitive offloading. *Trends in Cognitive Sciences*, 20:676–688, 2016. doi: 10.1016/j.tics.2016.07.002.

Steven Samuel, Geoff Cole, and Madeline J. Eacott. Grammatical gender and linguistic relativity: A systematic review. *Psychonomic Bulletin & Review*, 26:1767–1786, 2019. doi: 10.3758/s13423-019-01652-3.

Sister Miriam Joseph. *The Trivium: The Liberal Arts of Logic, Grammar, and Rhetoric*. Paul Dry Books, Philadelphia, 2002. First published 1937; revised edition used here.

Alexander Skulmowski. The cognitive architecture of digital externalization. *Educational Psychology Review*, 35, 2023. doi: 10.1007/s10648-023-09818-1.

Betsy Sparrow, Jenny Liu, and Daniel M. Wegner. Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333:776–778, 2011. doi: 10.1126/science.1207745.

Matthias Stadler, Maria Bannert, and Michael Sailer. Cognitive ease at a cost: LLMs reduce

mental effort but compromise depth in student scientific inquiry. *Computers in Human Behavior*, 160:108386, 2024. doi: 10.1016/j.chb.2024.108386.

Benjamin C. Storm, Sean M. Stone, and Aaron S. Benjamin. Using the Internet to access information inflates future use of the Internet to access other information. *Memory*, 25:717–723, 2017. doi: 10.1080/09658211.2016.1210171.

Shira Tal and Inbal Arnon. Redundancy can benefit learning: Evidence from word order and case marking. *Cognition*, 224:105055, 2022. doi: 10.1016/j.cognition.2022.105055.

Guillaume Thierry. Neurolinguistic relativity: How language flexes human perception and cognition. *Language Learning*, 66:690–713, 2016. doi: 10.1111/lang.12186.

Peter Trudgill. *Sociolinguistic Typology: Social Determinants of Linguistic Complexity*. Oxford University Press, Oxford, 2011.

Astrid Vereeck, Mark Janse, Katja De Herdt, Cedric Hauspie, Arnaud Szmalec, and Evy Woumans. Does studying Latin make pupils smarter? Presenting the field of classical language impact studies. *The Classical Journal*, 120:479–519, 2025. doi: 10.1353/tcj.2025.a959698.

Phillip Wolff and Kevin J. Holmes. Linguistic relativity. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2:253–265, 2011. doi: 10.1002/wcs.104.