

Toward AI-Scaffolded Pedagogical Stylistics: A Conceptual Framework for Bridging the Language–Literature Divide in L2 English

Yasir Hazim Qasim

Directorate of Education of Basrah Governorate, Ministry of Education, Iraq

Email: yasir.hazim.qasim@gmail.com

Received July 15, 2026; revised August 26, 2026; accepted August 27, 2026

Abstract

Literature has a long-standing place in second-language (L2) English teaching because it offers learners language that is rich, patterned, and affectively charged, and pedagogical stylistics provides a principled way of turning that language into learning by having students notice the choices through which literary effects are achieved. Recent classroom evidence is sobering. Large-scale observation and competence profiling show that literary texts still appear in L2 classrooms yet work with them seldom moves beyond literal comprehension and personal reaction; furthermore, learners' aesthetic-stylistic competence lags well behind their interpretive and empathic competence. The reason is structural rather than attitudinal: contingent stylistic guidance, responsive to what each reader notices, cannot be given to thirty readers at once. This article proposes AI-Scaffolded Pedagogical Stylistics (ASPS), a conceptual framework in which a generative language model is positioned not as an interpreter but as a semiotic tool within a teacher-designed activity, one that issues questions directing attention to stylistic features while the learner's aesthetic response and interpretive authority remain primary. Its contribution is threefold: it defines the construct of AI-scaffolded stylistic noticing by conditions that separate it from neighboring ideas; it specifies the mechanism by which repeated, progressively faded questioning is internalized as a transferable reading routine; and it advances the distinction between scaffolding and substitution as a falsifiable claim, namely the prediction that the same tool used to question rather than to answer will yield more learner-initiated stylistic noticing, and better delayed transfer to unaided reading, a difference traceable in a transcript. A preliminary coding framework and a set of testable propositions set out how the claim can be examined empirically.

Keywords: reader-response theory, stylistic noticing, aesthetic reading, generative artificial intelligence, second-language reading

Introduction

For over thirty years literature has been promoted in second-language (L2) English teaching because it offers language that is dense, patterned, and emotionally resonant, and invites interpretation rather than simple decoding (Carter & Long, 1991; Hall, 2005; Lazar, 1993). The discipline that has most convincingly delivered on that promise is pedagogical stylistics, which guides learners to read closely for the linguistic choices, the foregrounding, deviation, and figurative patterning through which literary texts make meaning (Hall, 2005; Short, 1996; Simpson, 2014). It works in two ways at once: it boosts interpretation and builds sustained attention to the target language into the substance of the lesson. But the old division between language teaching and literature teaching, shown in Paran's (2008) survey, was never completely settled, and the stylistic reading that would bridge that division is still patchy in classrooms.

Recent evidence makes the issue sharper. In a large video-based study of around 200 English lessons, Dodou and Brevik (2025) found that literary texts were present and read often but seldom explored in the layered, reflective way that justifies their use. Calafato and Hunstadbråten (2025) described the literary competence of upper-secondary EFL learners and revealed an instructive asymmetry: interpretive and empathic competence were relatively well developed, but aesthetic-stylistic and cultural-discursive competence, the ability to attend to how language is shaped and to locate a text within its discursive frame, lagged. It is an uncomfortable paradox. Of the various dimensions of literary reading, the stylistic dimension, which most directly serves language learning, is the least developed in current practice.

The deficit is structural. Bloemert et al. (2016) advocate a comprehensive strategy that attends to text, context, reader, and language, but classroom studies consistently show that teachers enact only the reader and language orientations, the latter often reduced to vocabulary glossing instead of stylistic analysis (Bloemert et al., 2019; Sun, 2023). Stylistic noticing is precisely what disappears when classes are large and time is short. The problem is not the value of reading stylistically but the difficulty of guiding thirty readers through it at once.

Generative artificial intelligence (AI) introduces a possibility here, though the literature gives little reason to expect it to be realized on its own. Large language models can sustain fluent, responsive exchange about a text with every learner at once (Kohnke et al., 2023), which is exactly the contingent attention a single teacher cannot supply. Yet research on AI in language education has concentrated overwhelmingly on written production: reviewing empirical studies from 2023 and 2024, Li et al. (2025) found that writing was the focus in 42.4% of the studies, with reading far behind and, within reading research, the literary, stylistic, and reader-response strands almost untouched (Teng, 2024; Xiao et al., 2026). The tool most able to scaffold stylistic reading has not been turned to that purpose.

This article asks how that capacity might be turned to stylistic reading. Used as an interpreter that answers learners' questions, a language model corrodes literary reading, dissolving the transaction between reader and text into the consumption of ready-made meaning. Used as a questioner that directs attention and returns the interpretive work to the learner, it can extend pedagogical stylistics to classrooms that cannot otherwise sustain it. The difference is a matter of design. This article is conceptual rather than empirical: it develops the framework, specifies its constructs and mechanism and derives testable propositions, leaving the empirical tests to future work. The article develops this design as AI-Scaffolded Pedagogical Stylistics (ASPS), a conceptual framework for the stylistic reading of literary texts in L2 English in which a generative model, treated throughout as a semiotic tool operating inside a teacher-designed activity rather than as a pedagogical agent, prompts stylistic noticing while the learner's aesthetic response remains central. The contribution is threefold and, taken together, more than a synthesis of its sources. First, the article defines a new construct, AI-scaffolded stylistic noticing, with conditions precise enough to separate it from ordinary noticing, from unaided close reading, and from its degenerate twin, AI-generated interpretation. Second, it outlines a process, dependent upon the stylistic fading at scale, by which an articulate model overcomes a structural impossibility that has hindered literature pedagogy for a century: one teacher cannot provide contingent, individually faded stylistic guidance to thirty readers in thirty different transactions. Third, it proposes the scaffolding-substitution distinction as a falsifiable design principle: the prediction that the same model used as questioner rather than as answerer will yield more learner-initiated stylistic noticing within the session and better delayed transfer to unaided reading, a difference whose interactional signature can be found in a transcript. The prediction is deliberately narrow. A supplied interpretation may well improve immediate comprehension; the claim concerns who does the interpretive work and what survives the session, not every measure that might be taken. The remainder maps out the problem, sets out the theoretical foundations, specifies the framework and its mechanism, weighs its risks and boundary conditions and proposes a research agenda by which each claim can be tested.

The Problem Space: A Persistent Divide and a Documented Deficit

An Unfinished Argument

The divide between language and literature teaching predates applied linguistics. Literary texts were central under grammar-translation, fell out of favour under audiolingual and early communicative approaches, and returned with the pedagogically oriented scholarship of the 1980s and 1990s (Carter & Long, 1991; Lazar, 1993). That revival rested on stylistics as the bridging discipline: by making the analysis of linguistic choices the route into interpretation, stylistic pedagogy promised to serve language and literature together (Short, 1996; Simpson, 2014; Widdowson, 1975). The promise

was articulate; its empirical verification was not. Paran (2008) showed that strong claims about literature's contribution to language development rested on a thin evidential base, and that integration was more aspiration than documented practice. The difficulty has never been the rationale, which remains compelling (Hall, 2005; Kramsch, 1993), but the gap between what literature is theorised to do and what classrooms are observed to do with it.

Model Versus Enactment

The most useful articulation of the integrationist ideal is provided by Bloemert and colleagues, who argue that literature is best taught by combining attention to the text as a crafted object, its context, the reader's response, and the language through which it works (Bloemert et al., 2016). The model works in principle; the empirical follow-up is sobering. Surveying experienced secondary EFL teachers in China, Sun (2023) found only the language and reader orientations practiced with any regularity, with text and context largely absent, and a similar contraction is reported in student-facing research (Bloemert et al., 2019). If teachers are not prepared and supported over time, the comprehensive strategy becomes partial, and the learners react personally, gloss vocabulary, and the stylistic core quietly disappears.

A Measurable Deficit in Stylistic Competence

The most telling evidence that this shortfall matters is that of Calafato and Hunstadbråten (2025). They evaluated literary competence in terms of interpretive, empathic, aesthetic-stylistic, and cultural-discursive dimensions. Upper-secondary EFL learners were significantly stronger on the first two than the last two. Learners could arrive at a reading and connect with characters, but were weak at perceiving and reasoning about the linguistic and formal choices that produce literary effects, and at situating texts within their discursive and cultural frames. The authors recommend more frequent close reading and analysis of techniques such as figurative language, symbolism, and narrative structure. With Dodou and Brevik's (2025) classroom observations pointing the same way, the convergence of an outcome study and an observation study marks a systematic feature of practice rather than an artefact of one cohort.

The Production Turn in AI, And the Double Gap

While pedagogy has under-served stylistic reading, AI scholarship has not corrected the balance, having been preoccupied with output (Li et al., 2025; Teng, 2024; Xiao et al., 2026). One study marks the current frontier and points to what lies beyond it. Lin et al. (2025) found AI-supported scaffolding to be more effective than teacher-led scaffolding in terms of comprehension, motivation, and attitudes. However, their texts were expository and their scaffolding was focused on comprehension and not perception of style or aesthetic response. Thus, two gaps coincide: pedagogy of

literature under-develops the aesthetic-stylistic competence it is best positioned to build, and research on AI is barely touching literary reading. A framework that uses AI to scaffold stylistic noticing would answer both, and the rest of this article develops one.

Theoretical Foundations

ASPS is based on the convergence of four established positions and one qualified claim about technology. Pedagogical stylistics provides the object of learning; Rosenblatt's transactional theory provides the constraint that protects the reader's experience; sociocultural theory provides the mechanism of support; the noticing hypothesis links attention to form with language development. The qualified claim concerns what a language model can and cannot be in this design.

Pedagogical Stylistics: Reading Language as Design

Stylistics is concerned with how linguistic choices create meaning and effect in texts (Leech & Short, 2007; Simpson, 2014). Pedagogical stylistics directs that enterprise towards learning, guiding learners to observe foregrounding, deviation, repetition, and figurative organisation so that interpretation deepens and language awareness sharpens together (Hall, 2005; Short, 1996). Its central wager is that interpretation grounded in observable textual evidence is more defensible and more instructive than interpretation asserted in the abstract. For the language learner, the by-product is consequential: attending to why a writer chose this word, this clause, or this image is a powerful form of focused engagement with the target language. Hanauer's (2001, 2012) work on the slow, attentive reading poetry demands shows that such attention is teachable and yields linguistic as well as interpretive gains. The competence pedagogical stylistics is built to develop is, in effect, the aesthetic-stylistic competence Calafato and Hunstadbråten (2025) found wanting.

The Transactional Theory of Reading

Rosenblatt's (1978, 1982) transactional theory supplies the indispensable counterweight to any analytic programme. A literary work, in this account, is not an object on the page but an event occurring in the transaction between a particular reader and a particular text. Rosenblatt distinguished two stances along a continuum: the efferent, concerned with information to carry away, and the aesthetic, concerned with the lived-through experience of the reading. Literary reading is predominantly aesthetic, and the danger this poses for stylistic pedagogy is well known: analysis can slide into an efferent hunt for technical labels, and the experience that made a text worth reading evaporates. Any framework that introduces a powerful analytic aid—and an articulate model is exactly that—must therefore treat the protection of the aesthetic stance as a design constraint rather than an afterthought. Stylistic noticing is to serve the transaction, never to replace it.

Scaffolding and the Zone of Proximal Development

Scaffolding, in the original formulation of Wood et al. (1976), is the support a more capable partner supplies so that a learner can accomplish what they could not manage alone, and it operates within what Vygotsky (1978) termed the zone of proximal development, the distance between what a learner can do unaided and what they can do with guidance. Two features of genuine scaffolding matter here, and both are drawn from that literature (van de Pol et al., 2010). The first is contingency: support responds to the learner's moment-to-moment performance rather than following a script. The second is fading: support is withdrawn as competence grows, so the learner internalises the activity. In these terms the structural problem of the literature classroom becomes clear. One teacher cannot give contingent, individually faded scaffolding to thirty readers in thirty separate transactions with a text; the stylistic guidance pedagogical stylistics needs is not deliverable by one person at the necessary grain.

Noticing and the Form–Meaning Link

Why does attention to form matter for acquisition? The bridge is provided by Schmidt's (1990) noticing hypothesis: features of the target language are unlikely to be acquired unless they are, at some level, consciously registered. A stylistic specialization of this principle is stylistic noticing. The learner attends not only to a grammatical form but also to its aesthetic function (how a passive defers an agent or a metaphor compresses two domains into one). Such noticing is doubly useful, both in deepening interpretation and in foregrounding target-language structure in an affectively memorable context, thus linking stylistics to second language acquisition.

The Large Language Model as a Semiotic Tool and Its Limits

The last component is the most delicate, and its theoretical status must be stated precisely. A language model can sustain fluent, responsive exchange about a text with every learner at once (Caines et al., 2023; Kohnke et al., 2023), which looks like an answer to the scaffolding problem, and the signal from adjacent work is positive (Lin et al., 2025). Yet the model is not a more knowledgeable other in the Vygotskian sense, and nothing in this framework depends on its being one: it does not understand the text, has no aesthetic experience, no communicative intent, and can state falsehoods with confidence. It is better described as a semiotic tool, a mediating artefact that carries and delivers questions designed by a teacher or materials writer, in the way a worksheet or a set of guided-reading cards does, but with the capacity to select its next question in response to what the learner has just written. Contingency, on this account, is a property of the designed interaction and of the prompt architecture that governs it, not a pedagogical capacity residing in the model. The distinction matters practically as well as theoretically. Treated as an authority, the tool becomes a dispenser of interpretations and a threat to both the transaction and the analysis; confined to the role of questioner, it can deliver the kind of contingent questioning

a skilled teacher provides, at a scale no teacher can match. ASPS depends on enforcing the second role and forbidding the first, and the design principles exist to make that enforcement explicit.

The ASPS Framework

Definition and First Principles

AI-Scaffolded Pedagogical Stylistics (ASPS) is an instructional framework for the stylistic reading of literary texts in L2 English in which a generative language model, used as a semiotic tool within a teacher-designed activity, delivers contingent prompts that direct the learner's attention to stylistic features and elicit the learner's own reasoning about their effects, while the learner's aesthetic response remains central and interpretive authority stays with the learner. Three commitments follow and govern every design choice: the AI prompts noticing rather than supplying interpretation; the aesthetic transaction is protected as the primary good; and support is contingent and designed to fade. These are formalized as the three principles in Table 1.

Table 1

The Three Design Principles of AI-Scaffolded Pedagogical Stylistics

Principle	Statement	Design Implication
P1. Scaffolding over substitution	The AI directs attention and elicits reasoning; it does not deliver finished interpretations.	Prompts are interrogative and evidence-seeking ("what do you notice about..."), never expository ("this symbolises..."). The learner produces the interpretation.
P2. Aesthetic primacy	The reader's lived, aesthetic transaction with the text is the primary good and is protected from efferent reduction.	Each cycle opens and closes with the reader's own response; analysis serves that response, and divergent personal readings are validated.
P3. Dialogic contingency and fading	Support responds to the learner's prior turn and is progressively withdrawn as stylistic noticing becomes independent.	Prompts adapt to learner responses; scaffolding intensity falls across a sequence so learners increasingly initiate noticing unprompted.

What is New: Three Contributions

It is worth stating plainly what ASPS adds, since its components are individually familiar. The novelty lies not in the parts but in their combination into a specified construct, a mechanism, and a testable design principle. The first contribution is the construct of AI-scaffolded stylistic noticing, defined in Section 4.5 by conditions that mark it off from neighbouring ideas: Schmidt's noticing concerns grammatical form for acquisition rather than aesthetic effect; stylistic close reading is neither dialogic nor faded to the individual reader; dialogic teaching is usually general rather than stylistic and, in most of its classroom forms, depends on the teacher's presence

(Alexander, 2020); intelligent tutoring pursues task correctness rather than literary response; and AI-generated interpretation hands the learner a product rather than a capacity. The construct occupies the position none of these fills (Table 2). The second contribution is the mechanism set out in Section 4.3, and the third is the scaffolding-substitution distinction as the principle deciding whether the same tool teaches reading or replaces it. Where this matters is worth spelling out. Research on technology in language learning has largely treated generative AI as an answer generator or a writing aid; the framework instead locates its value for reading in the questions it asks rather than in the interpretations it supplies. For literary pedagogy it offers a theoretical answer to an old structural constraint, the impossibility of giving contingent stylistic guidance to thirty readers at once. It also gives empirical research a clear target, since the scaffolding-substitution distinction marks the point at which the framework can be tested and, should the prediction fail, refuted.

Table 2

AI-Scaffolded Stylistic Noticing Distinguished from Adjacent Approaches

Approach or Construct	Object Of Attention	Who Produces the Interpretation	Contingent And Faded to Each Reader at Scale
Schmidt's noticing	A grammatical form, for acquisition	The learner registers the form	Not an interpretive act
Stylistic close reading	Linguistic choices and their effects	The learner, guided by the teacher	Rarely; one teacher can seldom fade support for every reader
Dialogic teaching	Meaning and reasoning, in general	The learner	Typically partial; usually human-mediated and rarely focused on style
One-to-one and intelligent tutoring	Task correctness and comprehension	The tutor or system	Yes, but not aesthetic or stylistic
AI-generated interpretation	The text's meaning, delivered	The model	Not applicable; unlikely to build interpretive capacity
AI-scaffolded stylistic noticing (ASPS)	A stylistic feature and its aesthetic effect	The learner, on cited evidence	Yes; contingent, faded, and individualized


Note. The columns describe typical rather than necessary features, and individual implementations vary. On contingency and fading as the defining properties of scaffolding see Wood et al. (1976) and van de Pol et al. (2010); on the range of classroom forms dialogic teaching takes see Alexander (2020); on what learners do with AI-supplied answers see Lin et al. (2025) and Xiao et al. (2026).

The Mechanism: How ASPS Produces its Outcomes

ASPS is not a bundle of theories but a causal claim about how a transferable stylistic-reading capacity is produced, and the chain has four links, rehearsed across successive cycles as support falls away (Figure 1). First, a contingent noticing prompt (P1) directs the learner to a specific feature without naming its effect. Second, the learner produces an observation and an evidence-grounded claim about effect, an instance of

AI-scaffolded stylistic noticing, which by the noticing hypothesis registers the form in an affectively charged context. Third, the cycle repeats across features while support fades (P3), so the routine of noticing, naming, and reasoning is rehearsed with decreasing assistance, which the framework predicts will in time be internalized as a strategy the learner can run alone. Fourth, throughout, aesthetic primacy (P2) keeps the felt transaction central and blocks the slide into efferent label-hunting that would void the gains. The proximal outcome, and the only one a transcript can show, is a developing stylistic-reading routine: the same noticing move performed with less help each time. Internalization and transfer lie beyond the session and are later outcomes to be tested separately, on an unseen comparable text after a delay, alongside aesthetic-stylistic competence and, through repeated form-meaning attention, incidental language gains. The moderator is dialogic role: a substitution-mode prompt short-circuits the chain at the second link, delivering a product the learner did not build, so no routine is internalized. The mechanism therefore predicts not that AI helps reading in general but that scaffolding-mode and substitution-mode uses of the same tool diverge on two specific outcomes, learner-initiated noticing within the session and delayed transfer beyond it, which is the difference the propositions are built to test.

Figure 1
Fading across Successive ASPS Cycles and the Scaffolding–Substitution Moderator



Note. Shading and border weight in the left-hand column indicate the level of assistance, which falls across cycles (P3) as learner initiative rises; P2 frames every cycle. The chain ends, inside the session, at a developing routine, which is what a transcript can show; internalisation and transfer sit outside the session box because they are later outcomes to be tested rather than results the dialogue establishes. The dashed path on the right shows substitution-mode use breaking the chain, so that no routine is built.

The Pedagogical Architecture: A Three-Phase Cycle

ASPS organises a reading encounter into three phases mirroring the established logic of pre-reading, during-reading, and post-reading work while re-specifying each around stylistic noticing and aesthetic primacy (Table 3). In the pre-reading phase, the tool activates the linguistic and schematic resources the learner will need and primes the aesthetic stance without pre-empting interpretation. The during-reading phase is the stylistic core: the tool issues contingent noticing prompts keyed to features the learner meets, asks the learner to describe each feature and hypothesise its effect, and withholds its own reading while requiring textual evidence. The post-reading phase moves outward, inviting the learner to consolidate an interpretation, connect form to felt response, and reflect critically on the activity, including a comparison between the learner's reading and any reading the model might offer. That comparison falls outside the scaffolded cycle rather than inside it. C4 governs the noticing phase, where a supplied interpretation would pre-empt the learner's own; once the learner's reading has been written down and fixed, the model's reading can no longer displace it and becomes an object of critique rather than a source of answers. This last turn is where critical AI literacy is cultivated.

Table 3
The Three-Phase ASPS Pedagogical Cycle

Phase	Learner Activity and Stylistic Focus	AI Scaffolding Function	Illustrative AI Prompt
Pre-reading	Activating relevant lexis, grammar, and experiential schema; entering the aesthetic stance.	Primes resources and stance without supplying interpretation.	"Before you read, think of a time a season matched your mood. Keep that in mind as you read."
During-reading	Noticing and naming foregrounding, deviation, figurative patterning, and syntactic choice; hypothesising effects from textual evidence.	Issues contingent, interrogative noticing prompts; withholds its own reading; requires evidence.	"The poem opens with a question rather than a statement. What does asking, instead of telling, do to your sense of the speaker?"
Post-reading	Consolidating interpretation; linking form to felt response; reflecting critically on the AI's role.	Elicits synthesis; then, once the learner's reading is fixed, offers a reading of its own for comparison and critique.	"Write your own reading first and keep it. Once it is written the noticing task is over: ask me for mine, then tell me where we differ and why your evidence supports yours."

Two exchanges from that transcript show the movement. Early on, with support at its most intense, the tool quotes the text and points to an exact spot without naming anything: it asks the learner to hold on to the word "stopped," look at the opening words of the passage, and say whether anything is odd about them as sentences. It is the learner who notes that the phrases are very short, and that one of them has no verb

at all, and the technical term follows rather than precedes the perception. Several turns later, without any prompting whatsoever, the same learner volunteers that a phrase about waters newly retired from the face of the earth sounds biblical, like the Flood, and asks if that is doing something too. No previous turn has prompted attention to that line. The initiative for noticing has moved to the reader and the tool's response does not clarify the allusion or offer a reading but instead returns the question and asks the learner to stay with their own reading. This move from tool-located evidence to learner-initiated perception is what fading means in this framework, and what a transcript makes visible.

The Central Construct, Specified

The central construct of the framework is AI-scaffolded stylistic noticing. The conditions are (C1) the attention is on a particular, locatable linguistic feature of a literary text, such as a foregrounded pattern, a deviation or a figurative or syntactic choice; (C2) the learner describes that feature in their own words, with metalanguage only supplied later; (C3) the learner makes a claim about the aesthetic effect of the feature that is based on cited textual evidence; and (C4) the contribution of the model is restricted to prompts that focus attention and elicit reasoning, withholding any ready-made interpretation for the duration of the noticing phase. C4 is bounded by that phase. The post-reading comparison described in Section 4.4 falls outside it, since a model reading solicited after the learner's own has been written can be criticised but can no longer pre-empt it. The conditions are diagnostic. If C4 fails, the episode is substitution, the degenerate twin in which the model explains, and the learner receives a product; if C3 fails it is mere labelling; if C1 fails it is generic discussion. Three further constructs can be defined as sharp. Aesthetic primacy is the requirement that the learner's felt response opens and closes the cycle and is never subordinated to analysis, enforced by having the learner commit to a personal reading before any analytic prompt. Contingent fading is the adjustment of prompt explicitness across a sequence in response to what the learner's previous turn shows they can do; it is not a schedule fixed in advance. Three observable states govern the next move. Support drops one level when the turn satisfies C1 to C3 with no help beyond the prompt already given, that is, when the learner locates a feature, describes it in their own words, and grounds a claim about effect in cited text. Support is held at the same level when the turn is partial, as when a feature is located and described but no effect is claimed, or an effect is claimed with nothing cited. Support rises one level when the turn misses the feature, asks the tool to supply the reading, or offers a label with no perception behind it. The levels are the four recorded in the coding framework of Section 6.1: quoting the text and naming the spot, a directional cue, an open invitation, and no prompt at all. The locus of interpretive agency is the identifiable origin, in a transcript, of each interpretive claim, learner or model; it is the variable that separates scaffolding from substitution and the manipulation check for Proposition 4.

Illustrations across Genres and Proficiency Levels

A single line isolates the mechanism. In the opening of Shakespeare's Sonnet 18, substitution mode answers the learner's question about what the line means and leaves no enlarged capacity, whereas an ASPS exchange asks what beginning with a question rather than a statement does to the reader's relation to the speaker, and has the learner predict whether the comparison will hold and read on to test it. The architecture extends across genres and levels. In narrative prose, a learner meeting the almost verbless opening of Dickens's *Bleak House* (1853) can be asked what the piled-up noun phrases do to the sense of time before any term such as parataxis is offered. In drama, a learner reading a curt exchange can be asked why a character answers a question with a question. Proficiency determines the grain, not the activity: a B1 learner meeting a refrain might be asked only whether the repeated line sounds the same the second time, with the term supplied once the difference has been felt, while a C1 learner is left to judge which feature is worth attention at all.

Calibrating Scaffolding to Proficiency

The transcript shows a B2 learner, and a fair concern is that the framework presumes a metalinguistic fluency that intermediate learners have yet to develop. The reply is that scaffolding intensity and prompt language are variables tuned to proficiency rather than constants: at B1 the tool works in the learner's simplest descriptive language, holds to one salient feature, and supplies metalanguage only after the phenomenon has been noticed; at B2 prompts can request effects directly; at C1 support fades fastest, withholding location cues so the learner decides what is stylistically salient. It is the same activity, just a different grain of aid. Proposition 6 is an empirical claim about the location of the lower boundary. To leave it open is more honest than to claim a universality not supported by the evidence.

Critical Considerations and Boundary Conditions

A framework that places a generative model between a learner and a literary text inherits the liabilities of that technology, and honesty requires naming these alongside the framework's own boundary conditions. Three kinds of limit need to be distinguished, since they are answerable in different ways: epistemological risks to the reading itself, socio-technical limits arising from the tool and its distribution, and pedagogical boundaries that mark where the framework ceases to apply.

Epistemological Risks

The most immediate is hallucination, in which a model confidently asserts that a text has a feature that it doesn't. The risk is partly contained because ASPS asks more than it tells, and any assertion must be based on verifiable textual evidence. But containment

is not elimination and the boundary condition is explicit: ASPS presupposes teacher mediation and a willingness to verify rather than uncritical acceptance. Interpretive homogenization is a second risk. If many readers put the same questions to the same model about the same canonical text, its statistically central reading may eclipse the divergent personal responses Rosenblatt holds essential; aesthetic primacy is the countermeasure, since learners must commit to their own reading before any model reading is solicited. A third is over-reliance and the deskilling it breeds, because a scaffold that never fades becomes a crutch, which is why contingency and fading must be built into the sequence rather than left to the learner's discretion.

Socio-Technical Limits

The risk of over-reliance is more acute, more technical. A model has a bias to be helpful, and a learner who simply demands, "just tell me what it means," will often find it obliging; where learners are motivated to extract answers, direct requests shade into deliberate attempts to talk the tool out of its role. If the ban on substitution exists only in the prompts the learner sees, it is one stubborn request away from collapse. Therefore, the principle must be enforced where the learner can't reach it, in the persistent system instructions governing the tool beneath the surface conversation, which fix its role, instruct it to answer demands for the answer by turning attention back to the text, and hold its function as questioner rather than oracle. A realistic position requires three caveats. Because models are probabilistic, system-level constraints increase reliability of behaviour without guaranteeing it. Sustained adversarial prompting can still elicit a finished interpretation, so the design reduces the risk rather than removing it, and teacher oversight remains inside the framework rather than outside it. The governance of these prompts, who writes them and how they are audited across model updates, becomes part of materials design. And the classroom routines described below carry as much of the burden as the technical settings do.

In classroom practice the boundary is held by a few concrete routines rather than by the model alone. The teacher or materials designer fixes the tool's role in a system prompt that learners do not edit, so the ban on giving answers does not rest on the wording each learner happens to see. Learners record their own reading before they may request the model's, which keeps the interpretive move on their side. And because the locus of interpretive agency is visible in a transcript, a teacher can read a sample of exchanges, judge quickly whether a learner is being prompted to reason or handed conclusions and reset the task when a dialogue has slid into substitution. Where a class is large, this need not mean reading everything: a short sample per learner, or the two summary indices described in Section 6.1, is enough to show which dialogues need attention.

Against these liabilities sits an opportunity. Bellot and Gutiérrez-Colón (2026) found that having learners compare AI- and peer-generated continuations of a literary text encouraged a critical stance toward the machine's authority and a sharper sense of the text as a created artefact. ASPS builds this into its post-reading phase, where the comparison between the learner's reading and the model becomes a designed occasion for critical AI literacy, so the framework develops two literacies at once, the stylistic and the critical-digital.

None of this removes the teacher; it redistributes the teacher's labour, freeing them from supplying contingent prompts to every learner and refocusing attention on orchestration, verification, and the cultivation of critical stance. The redistribution could look like a heavier burden, asking one teacher to be stylistician, prompt designer, and monitor at once; the answer is not to have teachers design from nothing. The prompts ASPS relies on can be drawn from a curated bank organised by stylistic focus and proficiency (Appendix B), so that preparation amounts to choosing a text, selecting two or three templates that suit its salient features, and pasting a fixed role instruction that need not be rewritten between lessons. This keeps the workflow within reach of teachers without advanced digital skills or spare preparation time and moves the design burden onto the materials designers and teacher educators who build and validate the bank once. The demand on teacher education is real either way: teachers need preparation in pedagogical stylistics and in the design and critique of AI prompts together, a need studies independently identify as the binding constraint on integration (Bobkina et al., 2021; Erdem Coşgun, 2025; Sun, 2023).

Pedagogical Boundaries

Limitations of the framework itself, beyond those of the technology, should be stated rather than managed away. The central construct lacks a validated measure. The coding scheme in Section 6.1 is a starting point rather than an instrument with established properties. In principle the distinction between scaffolding and substitution is clear, but in practice it can blur, because a learner who reproduces a half-remembered classroom reading can provide the surface of noticing without its substance, and coding will not separate every case cleanly. Another limit is imposed by cognitive load, most obviously at B1, where noticing and expressing an effect in the target language may exceed learners' attentional capacity. The calibration in Section 4.7 offers a partial answer, by limiting the focus to one feature and deferring the metalanguage, and a teacher may sensibly allow the learner's first language for the reasoning turn, while the textual evidence is in English, so that limited L2 production does not obscure stylistic perception. Proposition 6 is an empirical question of whether such code-switching supports or dilutes noticing, not a recommendation. The applicability of the framework outside the canonical, alphabetic, English-language examples presented here is also not established: the constructs are not obviously English-specific, and

adaptation to L2 Spanish, French, or Mandarin is a plausible extension, but it is proposed here as an empirical possibility to be tested rather than a demonstrated claim, as what counts as a salient stylistic feature and how metalanguage is introduced will differ by language and tradition. Finally, it is worth saying where the framework is the wrong tool. Where no teacher intervenes, where the model at hand is too weak to support disciplined questioning, where a text's power is mostly non-stylistic, or where the aim is efferent information-gathering rather than aesthetic reading, a simpler approach is better, and ASPS should not be forced on the task.

Resourcing raises a question of justice the framework cannot dodge. ASPS assumes a model capable of fluent, contingent exchange, and access is uneven and often gated behind subscriptions; a framework working only where a school can afford a premium model would deepen the inequity it claims to ease. Taking equity seriously means asking whether ASPS can run on free, open-source models within reach of under-resourced classrooms. The honest expectation is that it will run less well, since smaller models hallucinate more and sustain disciplined exchange less reliably, so any claim to democratic access must test how far such models carry the framework. Implementations must also protect learner data, hold to transparent norms of attribution, and use only literary texts free of permission constraints.

These conditions are not incidental to the setting the framework most has in view. In many EFL systems the ordinary literature lesson is a large class taught by one teacher with limited contact time and few materials, which is exactly the situation in which contingent stylistic guidance is hardest to give by hand. A framework whose purpose is to supply that guidance to many readers at once is therefore of most use where the constraint is tightest, provided the equity conditions noted above are met.

A Research Agenda: Testable Propositions

ASPS is conceptual, and its value is to be determined empirically. The propositions below translate it into hypotheses, each with a feasible design and an outcome construct, together forming a programme through which the framework can be confirmed, qualified, or refuted (Table 4). The framework is disconfirmed if scaffolding-mode and substitution-mode prompts produce equivalent rates of learner-initiated noticing and equivalent delayed transfer (Proposition 4), since the claim on which ASPS stands is the difference between teaching reading and replacing it.

Proposition 1 concerns the core promise: An ASPS condition will produce greater gains in aesthetic-stylistic competence than a literature-only condition and will match rather than surpass a skilled teacher-led stylistics condition delivered at ordinary class size. The second contrast is deliberately framed as equivalence. Nothing established here shows that a model reads better than a good teacher; the claim is only that contingent stylistic guidance can reach a whole class at once,

which one teacher cannot manage. A variant worth testing alongside it is whether teacher-plus-ASPS instruction improves on ordinary classroom practice. The design is a quasi-experiment with a validated literary-competence measure of the kind used by Calafato and Hunstadbråten (2025), with the teacher-led contrast analysed for equivalence rather than superiority. Proposition 2 addresses the central risk: ASPS will not depress, and may enhance, the aesthetic quality of reading even as analytic depth rises, testable with a stance measure along Rosenblatt's efferent-aesthetic continuum alongside think-aloud protocols. Proposition 3 concerns linguistic uptake: ASPS will lead to gains in incidental vocabulary in context, extending the effects of text-type and approach reported by Graham et al. (2024). Proposition 4 is the most stringent internal test of the framework. Prompt type is the condition compared, not a moderator: prompts that elicit learner reasoning are set against prompts that supply interpretation, and the outcomes expected to differ are learner-initiated stylistic noticing within the session and delayed transfer to an unseen comparable text. Discourse analysis of learner-tool transcripts, coded for the locus of interpretive agency, serves as the manipulation check, confirming that the two prompt conditions were realised as designed; it records how the prompt worked and is not itself a learning outcome. Proposition 5 concerns the critical-digital goal: A mixed-methods study can test whether a structured comparison of learner and AI readings predicts gains in critical AI literacy, using a critical-stance instrument and an analysis of learner reflections. The discussion of proficiency leads to Proposition 6: a minimum proficiency threshold conditions the effectiveness of the framework, below which the demands of the dialogue turn support into frustration. A comparison of B1, B2, and C1 cohorts under graded scaffolding, tracking cognitive load alongside competence, would locate it and identify the adaptations that lower it.

A Preliminary Coding Framework for Learner–Tool Transcripts

The decisive claims of this framework are visible only in interaction, so the propositions need a way of reading transcripts that another researcher could apply and check. What follows is a starting point rather than a validated instrument. The unit of analysis is the exchange pair: one tool turn together with the learner turn that answers it, since neither turn on its own shows who did the interpretive work. Each pair is then coded on the dimensions set out in Table 5. Three of these matter most. The locus of interpretive agency asks whether the substantive claim about effect originates with the learner or with the model. The level of assistance records how much a prompt gives away, and because it is ordinal, it doubles as the fading measure. Noticing status applies the four conditions from Section 4.5 to decide whether a pair counts as AI-scaffolded stylistic noticing, as substitution, or as neither. Two decision rules settle the awkward cases. Where a learner restates a claim the tool has already made, the pair is substitution whatever the order of speaking, because origin rather than

final speaker is what the construct turns on. Where a learner offers a plausible reading but cites nothing from the text, the pair is labelling, since C3 has failed. Reliability would be established in the usual way: two coders work independently on a common subset, agreement is reported for each dimension separately, and disagreements are settled against the written rules before the rest is coded.

Fading is read across pairs rather than within them. A sequence of exchanges yields a trajectory of assistance levels, and the framework predicts that this drifts downward while the share of learner-initiated pairs rises. Two figures summarise a session: the mean assistance level, and the proportion of pairs the learner began. The distinction between what happens in the dialogue and what survives it is worth holding onto. A transcript can show that a learner noticed something with a certain amount of help at a certain moment. Whether the routine has been internalised is a different claim altogether, and the only way to test it is to take the tool away and watch the learner read an unseen but comparable text after some weeks have passed. Studies bearing on Propositions 1–3 should report both and should not let strong performance in dialogue stand in for evidence of transfer.

Table 4

A Research Programme for ASPS: Propositions, Designs, and Constructs

No.	Proposition	Suggested Design	Construct / Measure
P1	ASPS produces greater gains in aesthetic-stylistic competence than a literature-only condition, and matches teacher-led stylistics at ordinary class size.	Three-condition quasi-experiment, pre/post; equivalence test for the teacher-led contrast.	Multidimensional literary-competence rubric.
P2	ASPS increases analytic depth without depressing aesthetic engagement.	Quasi-experiment plus think-aloud protocols.	Efferent-aesthetic stance measure; depth coding.
P3	ASPS yields incidental vocabulary-in-context gains.	Embedded pre/post vocabulary testing.	Vocabulary-in-context knowledge test.
P4	Scaffolding-mode prompts produce more learner-initiated noticing and greater delayed transfer than substitution-mode prompts.	Prompt-type comparison with delayed post-test; transcript discourse analysis.	Learner-initiated noticing rate; delayed transfer to an unseen text; locus of interpretive agency (coded) as manipulation check.
P5	Structured comparison of learner and AI readings predicts gains in critical AI literacy.	Mixed methods; reflection analysis.	Critical-stance instrument; qualitative reflection.
P6	A minimum proficiency threshold conditions effectiveness; graded scaffolding lowers it.	Cross-level comparison (B1/B2/C1) under graded scaffolding.	Competence gains; cognitive-load and frustration measures.

Table 5

A Preliminary Coding Framework for Learner–Tool Transcripts

Dimension	What it Records	Observable Indicators in the Pair	Values
Locus of interpretive agency	Who originates the substantive claim about effect.	Who first states the effect; whether the learner’s claim restates a model turn.	Learner-originated/ model-originated/ joint
Level of assistance	How much the prompt gives away, providing the fading measure.	Quoting the text and naming a spot; a directional cue; an open invitation; no prompt.	4 (high) to 1 (none; ordinal)
Stance	Whether the felt reading stays in view.	Reference to personal response; pursuit of labels without effect.	Aesthetic/mixed/ efferent
Noticing status	Whether the pair instantiates the construct.	C1 locatable feature; C2 learner’s own words; C3 cited evidence; C4 no supplied interpretation.	Noticing/labelling/ substitution/other
Initiation	Whether the learner raises a feature unprompted.	A feature introduced with no directing prompt in the prior turn.	Learner-initiated/ tool-initiated

Note. The unit of analysis is the exchange pair. Assistance is ordinal, so a session yields a fading trajectory; noticing status applies conditions C1–C4 from Section 4.5.

Two notes apply across the programme. Because model behaviour evolves quickly, studies should specify the model class used and pre-register prompt protocols so findings remain interpretable as the technology changes; model capability is itself best treated as a variable, with frontier and open-source systems compared to show how far effects depend on the underlying model, a question of direct equity significance. And because the framework’s claims are mechanistic, the analysis of learner-AI transcripts is not decoration but the evidentiary core, the only place the decisive scaffolding-versus-substitution distinction can be observed.

Conclusion

The argument has moved from a documented problem to a designed response. L2 English classrooms still widely use literature, but instruction develops interpretive and empathic competence more effectively than the aesthetic-stylistic competence that most directly serves language learning because the stylistic core of integrated literature pedagogy is hard to deliver to many readers at once.

Generative AI has the responsiveness the problem demands but has been aimed almost entirely at written production. ASPS redirects it, positioning a model as a semiotic tool that prompts stylistic noticing while interpretation and aesthetic experience remain with the learner, thereby delivering at scale the contingent guidance one teacher cannot. The deeper claim is modest: the recurring failure of literature-language integration has been a failure not of rationale but of delivery, and by recasting the model from a generator of interpretations into a scaffold for perception,

ASPS points toward recovering the stylistic orientation the comprehensive approach has most often lost.

Declarations

This research received no specific grant from any funding agency, and the author reports no conflict of interest.

Biodata

Yasir Hazim Qasim is an English language teacher and researcher at the Directorate of Education of Basrah Governorate, Ministry of Education, Basra, Iraq. His research interests include literature in language teaching, pedagogical stylistics, reader-response pedagogy, and artificial intelligence in English language education. <https://orcid.org/0009-0009-1299-7024>

References

- Alexander, R. (2020). *A dialogic teaching companion*. Routledge.
- Bellot, A.-R., & Gutiérrez-Colón, M. (2026). All assignments are equal, but some assignments are more equal than others: Exploring student responses to AI and peer-generated alternative endings in Orwell's *Animal Farm*. *Journal of Adolescent & Adult Literacy*, 69(4), Article e70025. <https://doi.org/10.1002/jaal.70025>
- Bloemert, J., Jansen, E., & van de Grift, W. (2016). Exploring EFL literature approaches in Dutch secondary education. *Language, Culture and Curriculum*, 29(2), 169–188. <https://doi.org/10.1080/07908318.2015.1136324>
- Bloemert, J., Paran, A., Jansen, E., & van de Grift, W. (2019). Students' perspective on the benefits of EFL literature education. *The Language Learning Journal*, 47(3), 371–384. <https://doi.org/10.1080/09571736.2017.1298149>
- Bobkina, J., Domínguez Romero, E., & Sastre-Merino, S. (2021). Literature and language education: Exploring teachers' views on teaching foreign language through literature in bilingual secondary schools in Madrid (Spain). *AILA Review*, 34(2), 145–186. <https://doi.org/10.1075/aila.21003.bob>
- Caines, A., Benedetto, L., Taslimipoor, S., Davis, C., Gao, Y., Andersen, Ø., Yuan, Z., Elliott, M., Moore, R., Bryant, C., Rei, M., Yannakoudakis, H., Mullooly, A., Nicholls, D., & Buttery, P. (2023). On the application of large language models for language teaching and assessment technology. *arXiv*. <https://arxiv.org/abs/2307.08393>
- Calafato, R., & Hunstadbråten, S. (2025). Literature in language education: Exploring EFL learners' literary competence profiles. *English Teaching & Learning*, 49, 537–559. <https://doi.org/10.1007/s42321-024-00193-w>
- Carter, R., & Long, M. N. (1991). *Teaching literature*. Longman.

- Dodou, K., & Brevik, L. M. (2025). From pages to practice: 10 years of teaching literature in English classrooms in Norwegian lower secondary school. *Reading Research Quarterly*, 60(4), Article e70039. <https://doi.org/10.1002/rrq.70039>
- Erdem Coşgun, G. (2025). Artificial intelligence literacy in assessment: Empowering pre-service teachers to design effective exam questions for language learning. *British Educational Research Journal*, 51(5), 2340–2357. <https://doi.org/10.1002/berj.4177>
- Graham, S., Zhang, P., Hofweber, J., Fisher, L., & Krüsemann, H. (2024). Literature and second language vocabulary learning: The role of text type and teaching approach. *The Modern Language Journal*, 108(3), 579–600. <https://doi.org/10.1111/modl.12946>
- Hall, G. (2005). *Literature in language education*. Palgrave Macmillan.
- Hanauer, D. I. (2001). The task of poetry reading and second language learning. *Applied Linguistics*, 22(3), 295–323. <https://doi.org/10.1093/applin/22.3.295>
- Hanauer, D. I. (2012). Meaningful literacy: Writing poetry in the language classroom. *Language Teaching*, 45(1), 105–115. <https://doi.org/10.1017/S0261444810000522>
- Kohnke, L., Moorhouse, B. L., & Zou, D. (2023). ChatGPT for language teaching and learning. *RELC Journal*, 54(2), 537–550. <https://doi.org/10.1177/00336882231162868>
- Kramsch, C. (1993). *Context and culture in language teaching*. Oxford University Press.
- Lazar, G. (1993). *Literature and language teaching: A guide for teachers and trainers*. Cambridge University Press.
- Leech, G., & Short, M. (2007). *Style in fiction: A linguistic introduction to English fictional prose* (2nd ed.). Pearson Longman.
- Li, B., Tan, Y. L., Wang, C., & Lowell, V. (2025). Two years of innovation: A systematic review of empirical generative AI research in language learning and teaching. *Computers & Education: Artificial Intelligence*, 9, Article 100445. <https://doi.org/10.1016/j.caeai.2025.100445>
- Lin, C.-C., Lin, T.-H., & Tang, C.-K. (2025). Enhancing English reading comprehension, learning motivation and attitude through AI-supported pre-reading scaffolding. *Journal of Computer Assisted Learning*, 41(6), Article e70150. <https://doi.org/10.1111/jcal.70150>
- Paran, A. (2008). The role of literature in instructed foreign language learning and teaching: An evidence-based survey. *Language Teaching*, 41(4), 465–496. <https://doi.org/10.1017/S026144480800520X>
- Rosenblatt, L. M. (1978). *The reader, the text, the poem: The transactional theory of the literary work*. Southern Illinois University Press.
- Rosenblatt, L. M. (1982). The literary transaction: Evocation and response. *Theory Into Practice*, 21(4), 268–277. <https://doi.org/10.1080/00405848209543018>

- Schmidt, R. (1990). The role of consciousness in second language learning. *Applied Linguistics*, 11(2), 129–158. <https://doi.org/10.1093/applin/11.2.129>
- Short, M. (1996). *Exploring the language of poems, plays and prose*. Longman.
- Simpson, P. (2014). *Stylistics: A resource book for students* (2nd ed.). Routledge.
- Sun, X. (2023). Literature in secondary EFL class: Case studies of four experienced teachers' reading programmes in China. *The Language Learning Journal*, 51(2), 145–160. <https://doi.org/10.1080/09571736.2021.1958905>
- Teng, M. F. (2024). A systematic review of ChatGPT for English as a foreign language writing: Opportunities, challenges, and recommendations. *International Journal of TESOL Studies*, 6(3), 36–57. <https://doi.org/10.58304/ijts.20240304>
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.
- van de Pol, J., Volman, M., & Beishuizen, J. (2010). Scaffolding in teacher–student interaction: A decade of research. *Educational Psychology Review*, 22(3), 271–296. <https://doi.org/10.1007/s10648-010-9127-6>
- Widdowson, H. G. (1975). *Stylistics and the teaching of literature*. Longman.
- Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2), 89–100. <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>
- Xiao, F., Zhu, S., & Xin, W. (2026). Exploring the landscape of generative AI (ChatGPT)-powered writing instruction in English as a foreign language education: A scoping review. *ECNU Review of Education*, 9(1), 1–19. <https://doi.org/10.1177/20965311241310881>

Appendix A. An Annotated ASPS Transcript

The following simulated transcript shows a B2 learner working with an ASPS tool on the opening of Charles Dickens's *Bleak House* (1853), a public-domain text. The left column gives the exchange; the right marks where each principle operates, with attention to the fading of support. The passage opens: "London ... Implacable November weather ... Fog everywhere. Fog up the river ... fog down the river ... Fog on the Essex marshes, fog on the Kentish heights."

Exchange	Design annotation
AI: Read the first two paragraphs once, slowly, without analysing. Then tell me in one sentence what this place feels like to you.	Aesthetic primacy (P2): the encounter opens with the learner's felt response, not analysis.
Learner: Heavy and grey, like everything has stopped. A bit suffocating, honestly.	A personal, aesthetic reaction precedes any technical work.
AI: Hold on to "stopped." Look at the opening words: "London. Implacable November weather. Fog everywhere." Anything odd about these as sentences?	Scaffolding over substitution (P1): the tool points to exact evidence and asks, without naming the feature. Support is at its most intense.
Learner: They're very short. And, wait, "Fog everywhere" has no verb. "Implacable November weather" isn't a full sentence either.	First noticing: the learner identifies the verbless construction, hedging.
AI: Exactly, verbless. Two of those phrases carry no verb at all. What does that do to your sense of time and movement?	Dialogic contingency (P3): the prompt builds on the discovery, pushing from feature to effect.
Learner: It makes it frozen. Nothing happens, things just are. That's why it felt stopped to me.	The learner links form to effect and back to the earlier felt response.
AI: You've tied the grammar to your first feeling. Now try the second paragraph: look at the word that opens one sentence after another, and tell me what the repetition does.	Support lightens to a directional cue plus an effect question; the tool no longer quotes the text.
Learner: It's "Fog": up the river, down the river, on the marshes. The repeating piles up; the fog feels everywhere and inescapable. It's overwhelming, on purpose.	Second noticing (anaphora) plus unaided reasoning; "on purpose" signals awareness of authorial design.
Learner: Actually, that line, "as if the waters had but newly retired from the face of the earth," sounds biblical, like the Flood. Is that doing something too?	Critical fading moment: the learner initiates noticing with no prompt, raising the allusion. Attention has passed from tool to reader.
AI: A sharp catch, and your own. Stay with your reading: why might Dickens reach for a world just after a great flood, here, on the first page?	The tool affirms and refuses to explain, returning the question; scaffolding is now minimal.
Learner: To make London feel ancient and primitive, civilisation slipped backwards, with all that mud and the earth half-formed. It fits the frozen, stuck feeling. I'd guess the fog comes back later as something nobody can see through.	Integrated interpretation plus anticipation of thematic resonance, unaided: evidence of transfer beyond the passage.
AI: Write that in your own words first. Then, if you like, ask how I'd read the fog, and tell me where we differ and whose reading the words better support.	Post-reading move enacting aesthetic primacy and critical AI literacy: the learner's reading precedes and is set against the model's, with evidence as arbiter.

The arc of support is the point. Early on the tool does the heavy lifting, quoting the text and pinning attention to a precise spot; through the middle it offers only directional cues; by the end the learner initiates both the noticing and the synthesis while the tool confirms and steps aside. That movement, from tool-led location to learner-initiated perception, is the fading the framework posits, and the practical signature that distinguishes genuine scaffolding from disguised substitution.

Appendix B. An Illustrative Bank of Prompt Templates

Templates are organised by stylistic focus and proficiency; they are starting points, not scripts, and the tool adapts each to the learner's previous turn in keeping with dialogic contingency (P3).

Stylistic focus	Generic template (bracketed slots filled from the text)	Proficiency calibration
Foregrounding & repetition	"You noticed that [feature] repeats. What happens to the feeling of the passage because the writer repeats it instead of varying it?"	B1: name the device only after the learner notices it. B2–C1: ask for the effect directly.
Figurative language	"The writer compares [X] to [Y]. What do the two share, and where does the comparison stop fitting?"	B1: focus on a single image. C1: invite judgement of how apt the comparison is.
Syntactic choice	"Look at the length and shape of these sentences. How does that change the pace or the weight of the scene?"	B1: contrast short versus long. B2+: tense, mood, clause order.
Narrative voice & point of view	"Whose eyes are we seeing through here? What does the narrator let you know, and what is kept from you?"	B2+: reliability of the narrator. C1: free indirect discourse.
Sound & rhythm	"Read the line aloud. Which sounds repeat, and what mood does that music create?"	Accessible at all levels; an effective entry point at B1.