



Designing AI-Mediated Arabic Learning Ecologies: Learner Agency, Target Activities, and Productive Constraints

تصميم بيئات تعلم اللغة العربية بوساطة الذكاء الاصطناعي: فاعلية المتعلم، والأنشطة المستهدفة، والقيود المنتجة

Aelita Miniianova

Russian islamic university, Russia

PhD student at Jordan university, Jordan

miniyanova@gmail.com

Abstract: Generative AI is reshaping Arabic education by redistributing linguistic work among learners, teachers, and technology. This conceptual-analytical study examines four adult learners across diverse Arabic-use contexts. Findings reveal that GenAI's value depends on alignment with learners' target activities and language ecologies, not support volume. The paper introduces the Target–Ecology–Gap–AI Affordance–Constraint–Agency heuristic. GenAI expands interaction and access to registers but should be constrained when it replaces target competence. Teachers should design learner–AI interaction while fostering learners' capacity for independent regulation

Keywords Arabic as a Foreign Language; generative AI; learner agency; pedagogical design; language-learning ecology

ملخص البحث

يُعيد الذكاء الاصطناعي التوليدي توزيع العمل اللغوي بين متعلمي العربية ومعلميها والتقنيات المستخدمة في تعليمها. تعتمد هذه الدراسة المفاهيمية التحليلية على أربع حالات لمتعلمين بالغين في سياقات متعددة للاستخدام اللغوي. وتُظهر النتائج أن فاعلية الذكاء الاصطناعي التوليدي إنما ترتفع بمدى تلاؤمه مع الأنشطة المستهدفة والبيئات اللغوية لكل متعلم، لا بحجم الدعم المقدم. وتستحدث الورقة الإطار الإجرائي: "الهدف-البيئة-الفجوة-الإمكانات الذكائية-المعوقات-الفاعلية الذاتية". فهو يعزز التفاعل والوصول إلى المستويات اللغوية المتخصصة، لكن يقيد حين يعطل تنمية الكفاءة المطلوبة. وعليه، ينبغي للمعلمين أن يصمموا سبل التفاعل بين المتعلم والذكاء الاصطناعي، مع إعداد المتعلمين لتنظيم هذا التفاعل باستقلالية

الكلمات المفتاحية: تعليم العربية للناطقين بغيرها؛ الذكاء الاصطناعي التوليدي؛ الفاعلية الذاتية للمتعلم؛ التصميم التعليمي؛ البيئة

اللغوية التعلمية

INTRODUCTION

Generative artificial intelligence has rapidly become part of the everyday ecology of language learning. Learners can obtain explanations, translations, corrective feedback, conversational practice, text revision, stylistic alternatives, and complete written texts almost instantaneously. AI is therefore better understood not as a set of isolated tools but as part of a wider pedagogical ecology. A systematic review of 221 studies in EFL/ESL education, for example, describes AI as occupying several roles at once: tutor, scaffold, conversational partner, assessor, coach, content generator, and support mechanism for teachers (Javadinejad & Davari, 2026).

This breadth of capability creates a distinctive problem for language education. Unlike many earlier technologies, generative AI can perform the target linguistic activity itself: it can write a message, formulate an argument, correct grammar, prepare a presentation, or sustain a conversation. Better task performance can no longer be assumed to demonstrate greater linguistic competence.

The central pedagogical question is therefore not simply whether, or how often, learners should use AI, but what AI should perform in relation to what the learner is expected to learn to do independently.

A 2026 scoping review of 123 studies on generative AI, cognitive offloading, and learner agency found two recurrent patterns. GenAI supported agency when learners exercised self-regulation, reflective engagement, feedback literacy, and active judgement; it weakened agency when they offloaded too much, accepted output uncritically, or replaced their own cognitive activity. The educational consequences thus depend less on the technology itself than on how it is embedded in learning (Wang et al., 2026).

Arabic language education makes this problem especially complex. Standard Arabic coexists with multiple spoken varieties, and learners may also need specialised academic, religious, media, administrative, diplomatic, or professional registers. A systematic review of 101 studies on Arabic diglossia confirms the continuing educational significance of diglossic distance and unequal exposure to Standard Arabic (Albirini et al., 2025).

Residence in an Arabic-speaking country does not necessarily provide sustained access to the forms of Arabic a learner needs to produce. Someone may manage everyday interaction in a local variety yet have few opportunities to formulate academic arguments in Standard Arabic. Another learner may read advanced specialised literature but rarely discuss it. Conversely, a highly proficient user living outside the region may retain strong receptive competence while gradually losing spontaneity in production.

The challenge is not simply to provide “more Arabic”, but to create access to the forms of linguistic participation required by a learner’s academic, professional, or personal goals.

This paper examines how generative AI can help create such access without taking over the competence learners seek to develop. It approaches GenAI integration from two connected perspectives: learner agency and pedagogical design. The latter entails identifying the target activity, mapping the learner's existing language ecology, locating the gap between them, selecting useful AI affordances, and constraining those affordances when they become substitutive.

The guiding research question is: How can teachers design AI-mediated Arabic learning environments that address gaps between learners' available language ecologies and their target linguistic activities while preserving learner agency?

Although grounded in Arabic language education, the design problem extends to multilingual and higher-education contexts in which everyday language use differs from specialised academic or professional participation.

1. METHODS

Theoretical Framework

1.1. Learner Agency Beyond Independence

A learner-centred orientation is relevant to AI-mediated language education, but it is too broad to explain what happens when learners can delegate substantial portions of linguistic activity to a generative system. Learner agency provides a more precise lens.

Agency does not mean working without assistance. Ecological accounts treat it as situated: learners act in relation to the goals, resources, constraints, people, and affordances available in a particular environment. Liu and Chao (2018) showed that technological affordances, teachers' pedagogical decisions, and learner agency operate interdependently in a technology-mediated language classroom.

Generative AI makes that interdependence more consequential because assistance can extend to producing the linguistic artefact itself. In an exploratory sequential study—interviews with 17 language learners followed by a survey of 646—Lai et al. (2026) identified several ways in which learners negotiate agency during self-directed, informal GenAI-supported language learning. Their findings frame learner–AI interaction as an ongoing negotiation of control and responsibility rather than a simple choice between autonomy and assistance.

In this paper, learner agency comprises the capacity to:

identify a meaningful target for language use;

make informed decisions about when and how AI should participate;

evaluate rather than merely accept AI-generated output;

retain control over linguistic actions central to the competence being developed; and

progressively shape the conditions of their own language learning.

Agency in this sense is not technological abstinence. A learner who uses AI strategically to create sustained opportunities for Arabic interaction may exercise more agency than one who avoids AI but has little meaningful contact with the language.

Agency is weakened, however, when AI performs the very activity the learner intends to acquire. The distinction between assistance and substitution is therefore central.

1.2. Target Activity and Needs Analysis

The second component of the framework comes from language needs analysis. Long (2005) argued that learner groups differ in the varieties, skills, genres, registers, and lexicons they require; generic programmes are therefore particularly inadequate for adults with identifiable academic or occupational purposes. Needs analysis begins not with a predetermined syllabus but with the activities learners need to perform through language.

This principle is especially relevant to intermediate and advanced learners of Arabic because broad labels such as B1, B2, or C1 can conceal substantial asymmetries.

One learner may understand specialised texts but be unable to discuss them spontaneously; another may speak fluently while making elementary grammatical errors in writing. An experienced Arabic teacher may not need to reach a higher CEFR level at all, but to prevent productive competence from weakening after leaving an Arabic-speaking environment.

The pedagogically useful question is therefore not only “What is the learner’s level?” but “What does this learner need to be able to do through Arabic?”

This orientation is consistent with recent work on Arabic for Specific Purposes, which emphasises needs-responsive curriculum design and the relationship between instruction and learners’ specialised communicative contexts (Salsabila, Fikroh, & Bazighoh, 2026).

Here, target activity refers to a meaningful linguistic activity that the learner aims to perform beyond the immediate exercise: teaching a disciplinary topic, writing a professional message, discussing an academic argument, managing everyday interaction, or maintaining sustained productive use of Arabic.

Such an activity may remain inaccessible even when the learner possesses substantial linguistic knowledge. What is missing may be not another explanation or text, but repeated participation in the activity itself.

1.3. Language-Learning Ecology and Affordances

An ecological perspective adds another question: what opportunities for performing the target activity are actually available to the learner?

The concept of affordance directs attention away from resources as objects and towards possibilities for action. A learner may be surrounded by Arabic yet still have few affordances for the particular kind of Arabic required.

The discrepancy is especially visible in a diglossic language. Living in Cairo, for example, may provide extensive everyday interaction in Egyptian Arabic, while university study offers lectures and reading in formal or specialised Arabic. Neither setting automatically provides repeated opportunities to formulate an academic position, defend an argument, teach a difficult concept in Standard Arabic, or receive focused feedback on academic writing. Physical immersion and target-relevant linguistic participation are not equivalent.

GenAI can alter this ecology by adding affordances that are otherwise absent or scarce. It may serve as an always-available interlocutor, a simulated student, an academic discussant, a feedback provider, or a source of alternative formulations.

Yet a capability is not pedagogically valuable simply because the system offers it. A polished AI-generated paragraph may be useful when the goal is discourse analysis and counterproductive when the goal is independent writing. AI-mediated pedagogy must therefore consider affordances together with constraints.

1.4. Teacher Agency and the Design of Learner–AI Interaction

Research on teacher agency shows that teachers do more than adopt or reject generative tools: they diagnose problems, select affordances, modify output, manage constraints, and exercise professional judgement.

In Jia and Pun's (2025) multiple-case study of teachers using ChatGPT, agency involved recognising language-related pedagogical problems, mobilising AI affordances, and mitigating technological and institutional constraints.

Morton and Peeters (2026) similarly describe AI-mediated language teaching as orchestration. Their study of tertiary English-language teachers identified iterative cycles of prompting, evaluating, and revising, as well as a teacher-as-editor role through which pedagogical judgement remained with the teacher.

Generative AI thus adds a new layer of pedagogical work. Alongside content, tasks, sequencing, interaction, and feedback, the teacher must design key features of the learner's interaction with an intelligent system:

what AI is expected to do;

what kind of support it may provide;

when it should intervene;

what it should deliberately withhold; and

how responsibility should gradually shift from teacher regulation to learner regulation.

The object of design is not merely a prompt but an interaction architecture. This becomes crucial when the technically most efficient use of AI is pedagogically undesirable.

2. Method and Analytical Approach

This study uses a practice-based conceptual-analytical approach. It neither reports an experimental intervention nor seeks to establish causal effects of GenAI on language proficiency. Instead, four contrasting cases from adult Arabic teaching and learning are used to examine how different target activities call for different configurations of learner–AI interaction.

The first is a reflective self-case involving an experienced Arabic-language educator; the other three derive from work with adult learners whose academic and professional profiles differ substantially. The cases are presented as pedagogical vignettes rather than a formal multiple-case study, and no claims of measured learning gains are made. Their analytical value lies in the contrasts among them.

Each case is examined through six questions:

What is the learner’s target linguistic activity?

What language-learning ecology is currently available?

What target-relevant affordance is missing or insufficient?

What additional affordance can GenAI provide?

Which GenAI capability should be constrained because it risks replacing the target activity?

What linguistic action or judgement must remain under learner control?

Together, these questions yield the practical design sequence developed in the cross-case analysis: Target → Ecology → Gap → AI Affordance → Constraint → Agency.

3. PEDAGOGICAL CASES

3.1. Advanced Proficiency Outside the Arabic-Speaking Environment: Embedding Arabic in an Existing Life Practice

The first case concerns an experienced Arabic-language educator with advanced proficiency who now lives outside an Arabic-speaking environment. The difficulty is neither limited grammatical knowledge nor an inability to understand sophisticated Arabic. Professional reading and occasional use continue, but sustained spontaneous production has become less frequent.

This pattern is common among advanced adult users. Receptive contact can be maintained through books, media, academic materials, and online content; productive interaction is harder to organise because interlocutors must be available, schedules must align, and the exchange must have enough substance to continue over time. Advice to “speak Arabic more often” identifies the desired behaviour without addressing these conditions.

The response is to embed Arabic in an activity that already occurs every day. The learner uses a detailed ChatGPT-based assistant to track nutrition, body weight, physical activity, and related reflections. Rather than creating a separate language-learning chatbot, the existing assistant is redesigned to conduct the entire interaction in Arabic.

The assistant maintains continuity across entries, responds to their substance, asks relevant follow-up questions, and notices recurrent errors without turning every exchange into a correction exercise.

Arabic thereby becomes the medium of an activity with a purpose beyond language learning, not the object of an exercise called “talking about health”.

Repeated entries elicit reporting, quantification, comparison, causal explanation, uncertainty, evaluation, planning, and reflection. Their value lies less in AI-generated instruction than in embedded continuity: GenAI lowers the organisational cost of sustained, meaningful production.

At this proficiency level, additional scaffolding may be unnecessary; the main design problem is to create a stable reason to produce Arabic every day. Because the content originates in lived experience, it remains learner-generated. AI sustains continuity, interaction, and selective feedback without composing the learner’s contribution.

This case suggests that GenAI is especially useful when it makes the target language the medium of an activity that already has independent value for the learner. Its pedagogical function lies not in inventing more exercises, but in reconfiguring an activity that already matters.

3.2. A B1+/B2 Graduate Student: From Receptive Exposure to Everyday Spontaneous Production

The second case also uses an AI assistant organised around an existing life activity, but the pedagogical design differs.

The learner is a graduate student who works at a university and has previous translation experience. She reads Arabic and regularly watches Arabic-language

videos, yet her productive everyday competence is considerably weaker than her receptive competence.

She has little regular speaking practice, experiences anxiety during spontaneous production, and has an uneven active vocabulary. She can handle relatively formal or familiar material but lacks parts of the everyday, emotionally expressive, colloquial-adjacent repertoire needed to discuss food, purchases, personal decisions, feelings, or disagreements naturally.

At the same time, she is managing her finances independently for the first time and genuinely needs to monitor income, expenditure, priorities, and spending decisions. A personalised financial assistant can support that non-language goal while creating recurring occasions for Arabic production.

Unlike the advanced educator, however, this learner needs graded communicative scaffolding. The teacher can structure recurring financial exchanges so that they elicit progressively more demanding discourse functions.

The sequence may progress from simple reporting to explanation, comparison, prioritisation, and reflective evaluation:

Reporting:

ماذا أنفقتِ اليوم؟

What did you spend today?

Explanation:

لماذا كان هذا المصروف ضرورياً؟

Why was this expense necessary?

Comparison:

هل أنفقتِ هذا الأسبوع أكثر أم أقل من الأسبوع الماضي؟ ولماذا؟

Did you spend more or less this week than last week, and why?

Prioritisation:

إذا كان لديك مبلغ محدود حتى نهاية الأسبوع، فما أولوياتك؟

If you have a limited amount of money until the end of the week, what are your priorities?

Reflective evaluation:

هل هناك قرار مالي كنت ستأخذينه الآن بطريقة مختلفة؟

Is there a financial decision you would now make differently?

Across these exchanges, the learner moves from reporting to explaining, comparing, justifying, evaluating, and planning. The central design decision is not the topic itself, but the degree and form of support.

When she struggles, the system should not immediately generate a complete answer for her to copy or repeat. It can instead ask a clarifying question, supply a lexical item, offer the beginning of a sentence, or provide another minimal cue that enables her to continue.

Constraint here is part of the scaffolding: generating a fluent answer may be the system's strongest technical response but its weakest pedagogical one when spontaneous production is the target.

AI therefore supplies both a scarce interactional affordance and adjustable support within the interaction. For the advanced user, the main need was continuity; for the B1+/B2 learner, continuity must be combined with lexical and interactional scaffolding. The tool alone does not determine its pedagogical role; that role emerges from the learner's profile, target activity, and existing ecology.

3.3. An Al-Azhar Graduate Student: Seven Years in an Arabic-Speaking Country Without Sufficient Target Practice

The third case shows why an ecological perspective is especially important in Arabic education. The learner is a graduate student in Sharia-related studies at Al-Azhar who has lived in an Arabic-speaking country for approximately seven years, yet his productive Arabic remains around B1+.

The profile appears paradoxical only until his actual language ecology is considered. He manages everyday communication, regularly reads specialised Arabic literature, understands substantial disciplinary content, and encounters formal Arabic in his studies. His future goals, however, require participation qualitatively different from receptive comprehension and everyday interaction.

His two principal target activities are teaching Sharia-related subjects and writing a graduate thesis in Arabic. Both require him to explain complex concepts, reformulate information for an audience, discuss specialised issues with professors, defend interpretations, construct arguments, and produce coherent academic prose.

His environment provides extensive Arabic input but far fewer opportunities for repeated academic production. A learner may understand a professor without responding with comparable conceptual precision, or follow a specialised text without reconstructing its argument independently. Familiarity with academic phrases and structures does not ensure that their productive use has become automatic. GenAI can help by creating several complementary academic micro-environments.

AI as an academic interlocutor

After reading an authentic disciplinary passage, the learner discusses it with the system. Rather than summarising the passage for him, AI asks questions such as:

ما القضية الرئيسية التي يناقشها الباحث؟

What is the main issue discussed by the researcher?

كيف يبرهن على رأيه؟

How does the author support this position?

هل توافقه؟ ولماذا؟

Do you agree, and why?

ما الاعتراض الذي يمكن أن يوجّه إلى هذا الرأي؟

What objection might be raised against this position?

AI does not replace comprehension here; it turns receptive contact into productive intellectual interaction.

AI as a student

For teaching practice, the learner assumes the expert role and AI becomes the student. The system can misunderstand a concept, ask for an example, request a simpler explanation, or challenge an argument.

Teaching requires more than possessing disciplinary knowledge: it requires restructuring that knowledge linguistically for someone else. The learner must therefore perform the sequence himself: explain → encounter misunderstanding → reformulate → exemplify → respond.

AI as a constrained academic reader

Academic writing requires a different design. If the learner submits an imperfect paragraph and asks AI to rewrite it in polished academic Arabic, the text may improve dramatically while his productive competence changes little.

A more appropriate sequence begins with oral formulation of the position, followed by an independently written paragraph and only then limited diagnostic feedback. AI might identify one register problem, one cohesion problem, and one unclear element in the argument without supplying a replacement paragraph.

The learner revises the paragraph himself. Only later may the system

generate an alternative for comparison; the learner then evaluates the two versions and explains what changed. This architecture recasts AI from a substitute author into a responsive academic environment.

The case illustrates a broader principle: immersion is not a property of geographical location alone. Pedagogically relevant immersion depends on access to the forms of participation required by the target activity. Arabic diglossia makes the discrepancy particularly visible, but the problem is not unique to Arabic.

3.4. A Professional Translator: When Good Pedagogy Requires Restricting AI

The fourth case reverses the logic of the preceding examples. The learner is a professional translator with many years of Arabic study, substantial experience in Arab countries, and strong oral communicative competence. He uses idioms and proverbs, speaks publicly, works with religious texts, and manages sophisticated spoken interaction.

His written Arabic is markedly less secure. Persistent problems occur in areas associated with foundational grammatical control, including constructions with *إِنَّ* وأخواتها, adjective agreement, the *idāfa* construction, verbal mood, and especially numerals.

The discrepancy is pronounced enough to undermine his confidence even in short WhatsApp messages, so digital tools frequently produce or check his texts. Yet his goal is explicit: to write correctly himself and understand why a form is correct.

That goal changes the pedagogical status of AI generation. Asking AI to compose a correct message is rational for professional efficiency, but as a learning strategy it removes the operation he needs to master. This case therefore requires deliberate restriction of the generative affordance.

One possible sequence is: independent first draft → AI diagnosis without rewriting → learner explanation of the relevant rule → learner revision → AI alternatives → learner comparison → final learner decision.

For a numeral error, for example, AI can identify the problem as one of numeral–noun agreement and require the learner to retrieve the rule rather than immediately supply the correct form. For a stylistic problem, it can offer alternatives only after an independent attempt and ask the learner to explain which version best fits the intended register.

What remains with the learner is not only the first attempt but the final linguistic judgement. Good AI pedagogy may therefore require disabling the system's most efficient capability: technological possibility is not, by itself, a pedagogical argument.

4. CROSS-CASE ANALYSIS

Across the four cases, GenAI has no fixed pedagogical function. An always-available conversational system creates continuity for one learner, scaffolds spontaneous production for another, simulates scarce academic participation for a third, and must be restricted for a fourth.

The relevant unit of pedagogical decision is not the AI tool itself, but the relationship among learner, target activity, and available ecology. The comparison yields a six-component design sequence:

Learner profile	Target activity	Ecological gap	Appropriate AI role
Activity retained by learner			

Advanced Arabic educator outside an Arabic-speaking environment	
Sustained productive Arabic use	Lack of regular meaningful interaction
Continuous personally relevant interlocutor	Formulating and reflecting on lived experience

B1+/B2 graduate student	Spontaneous everyday production	Limited speaking practice and active everyday vocabulary
adaptive scaffolding	Interlocutor with minimal Producing the message and making communicative decisions	

B1+ Al-Azhar graduate student	Teaching and academic Arabic	Limited productive academic participation
Academic interlocutor, simulated student, diagnostic reader	Explanation, argumentation, initial academic production	

Professional translator	Independent accurate writing	Grammatical instability and reliance on external correction
metalinguistic prompting	Diagnosis, comparison, First draft, revision decisions, final linguistic judgement	

Target → Ecology → Gap → AI Affordance → Constraint → Agency

Target. What does this learner ultimately need to be able to do through Arabic?

Ecology. What opportunities for that activity already exist in the learner’s current linguistic environment?

Gap. Which target-relevant forms of participation are absent, insufficient, or difficult to sustain?

AI affordance. What opportunity can GenAI add to the existing ecology?

Constraint. Which AI capability should be limited because it would replace rather than support the target activity?

Agency. What linguistic action, judgement, or responsibility must remain with the learner?

The sequence is a practical analytical heuristic, not a validated theoretical model. Its order matters: beginning with “What can ChatGPT do?” makes technological capability the organising principle of pedagogy; beginning with “What does this learner need to learn to do?” makes technology conditional on prior pedagogical diagnosis.

The framework therefore shifts attention from AI integration to AI allocation: which parts of a linguistic activity should be performed by the learner, which may be supported by AI, and which should not yet be delegated. As generative systems become more capable, capability itself becomes an increasingly poor criterion of pedagogical value.

5. DISCUSSION

5.1. Personalisation Beyond Adaptive Difficulty

GenAI is often associated with personalised learning because it can vary text difficulty, generate individual exercises, adapt explanations, and provide immediate feedback. Research with Arabic instructors similarly reports its use for materials, assessment, and personalised learning plans, alongside concerns about dialect processing, cultural authenticity, and assessment reliability (“Generative AI Implementation and Assessment,” 2025).

The cases examined here point to a more consequential form of personalisation. It begins not with a customised exercise, but with identifying why a particular person needs Arabic and which forms of linguistic action are missing from that person’s life.

The advanced educator needs an architecture that restores sustained productive use, not easier or harder materials. The graduate student needs repeated, low-risk opportunities for spontaneous everyday production and active-vocabulary growth. The Al-Azhar student needs productive academic participation; the translator needs grammatical control over independently written text.

In each case, the environment is personalised because its pedagogical function changes, not merely its content. Material-level customisation remains superficial if every learner is still asked to perform essentially the same linguistic activity.

5.2. Reconsidering the Meaning of Language Environment in Arabic

The Al-Azhar case also complicates the familiar geographical distinction between learning inside and outside a target-language environment.

For Arabic, geography alone is a poor guide. An Arabic-speaking city is linguistically rich, but that richness may not align with the learner’s target register or activity. A learner can have abundant dialectal interaction but little productive Standard Arabic, receive sophisticated academic input while remaining primarily a listener and reader, or need diplomatic, professional, or disciplinary writing that rarely appears in everyday interaction.

The more useful distinction is therefore between an available environment and a target-relevant environment.

This reframing matters for digital pedagogy. GenAI need not be treated as a synthetic substitute for “real” linguistic life; it can add a missing layer of participation to an existing environment. The Al-Azhar learner needs AI not because Arabic is absent from his surroundings, but because repeated academic production is insufficiently available.

Diglossia and register differentiation make the issue especially visible in Arabic, but it also arises in multilingual higher education whenever the language of everyday interaction differs from the language of academic or professional participation.

5.3. From AI Affordances to Productive Constraints

Technology-enhanced language pedagogy has traditionally asked what digital tools make possible. Generative AI requires a second question: which possibilities should remain unavailable during learning?

The answer depends on the target activity. An AI-generated first draft removes the central learning opportunity when the target is independent professional writing, but may be useful when the target is discourse analysis. A generated answer prevents spontaneous production, yet several alternatives may be exactly what a learner needs to notice stylistic variation. The pedagogical value of an AI operation is relational, not intrinsic.

This argument accords with Wang et al. (2026), who found that GenAI-supported learning cannot be reduced to a beneficial/harmful binary: scaffolded, augmentation-oriented practices are more compatible with agency, whereas replacement, overreliance, and uncritical uptake remain recurrent risks.

The present analysis goes further by treating constraint as a positive pedagogical resource, not merely a response to technological risk. A productive constraint preserves the difficulty through which learning occurs. Unlike a restrictive constraint imposed for external reasons, it limits technology because unrestricted assistance would remove the learner’s opportunity to perform a target linguistic operation.

5.4. Task Performance Is Not Competence

The cases also expose a problem beyond Arabic education: generative AI weakens the connection between successful performance and demonstrated competence.

A learner may submit a grammatically sophisticated paragraph without being able to reproduce its structure, a professional may send an idiomatic message without having formulated it, and a student may present an AI-generated academic argument that exceeds their ability to defend it. The quality of the artefact is therefore an increasingly unreliable indicator of what the learner can do.

Independent subsequent performance becomes a crucial criterion of evaluation. The question is not only whether AI helped the learner complete a task better, but whether, after using it, the learner can perform more of the relevant linguistic activity without support.

Research and classroom assessment both need to distinguish more clearly between AI-assisted performance and the development of independent capacity.

5.5. The Teacher as Designer of an AI-Mediated Language Ecology

The four cases complicate predictions that increasingly capable generative systems will reduce the importance of the language teacher. Some familiar functions can certainly be redistributed: learners no longer need a teacher whenever they want a grammatical explanation, example sentence, basic correction, or conversational partner.

That redistribution creates a different professional demand. The teacher works increasingly at a meta-pedagogical level: diagnosing asymmetric competence, identifying target activities, mapping the learner's ecology, selecting an AI affordance for a specific gap, calibrating scaffolding, constraining substitutive assistance, and checking that meaningful linguistic judgement remains with the learner.

This role is consistent with empirical work on teacher agency. Jia and Pun (2025) describe teachers as mobilising AI affordances while mitigating constraints; Morton and Peeters (2026) describe task-sensitive orchestration as continuing evaluation and revision rather than simple adoption.

For language education, the emerging role can be described as teacher as designer and orchestrator of learner–language–AI interaction. This does not require every teacher to become an AI specialist. Prompt construction is only the technical expression of an earlier pedagogical decision, and a sophisticated prompt built on a poorly diagnosed need remains poor pedagogy.

5.6. From Teacher-Regulated to Learner-Regulated AI Use

Placing the teacher at the centre of AI-mediated design may appear to conflict with learner agency. If teachers permanently decide how, when, and why AI may be used, learners remain dependent on external regulation. Effective design must therefore include the conditions for its own gradual withdrawal.

A plausible trajectory is teacher-regulated → co-regulated → learner-regulated AI use. Initially, the teacher may specify the interaction architecture: calibrating support for the B1+/B2 graduate student, withholding AI first drafts from the translator, or structuring the Al-Azhar student's academic dialogue.

At the next stage, the rationale for these decisions becomes explicit. Eventually, learners should be able to ask:

What exactly am I trying to develop?

What form of practice is currently missing?

What could AI usefully contribute?

What should I still do myself?

How will I evaluate its response?

Am I using AI to learn the activity, or merely to complete it?

Once learners can answer these questions, they begin to design their own AI-mediated language ecologies. Learner agency is not the opposite of pedagogical design; carefully designed external regulation can be the route to more independent regulation.

6. PEDAGOGICAL IMPLICATIONS

The analysis has five practical implications.

First, teachers of intermediate and advanced Arabic should replace broad diagnoses such as “needs more speaking” or “needs better writing” with analyses of specific target activities. “Speaking” may mean explaining a theological concept, negotiating a purchase, disagreeing politely with a colleague, or sustaining informal conversation; “writing” may mean producing an institutional email, constructing an academic argument, or composing a short professional message without external correction.

Second, language-ecology analysis should become part of needs analysis. Teachers need to ask not only what learners know, but where, with whom, in which variety and register, and for what purposes they currently use Arabic.

Third, prompts should be treated as components of instructional design rather than isolated technical devices. A prompt operationalises an earlier decision about the role AI should play.

Fourth, productive constraints should be designed deliberately. Useful instructions may include:

Do not rewrite the text.

Do not provide the answer immediately.

Ask one question first.

Give only one lexical cue.

Identify the type of error but require the learner to correct it.

Provide an alternative only after the learner has produced an independent version.

Fifth, personally meaningful activities beyond conventional language learning deserve more pedagogical attention. A health tracker, financial diary, professional planning system, academic reading routine, or research workflow can become a language-use environment when Arabic is the medium through which the activity is conducted.

This approach is particularly valuable for adults whose difficulty lies not in accessing materials, but in integrating sustained target-language participation into complex professional and personal lives.

7. LIMITATIONS AND FUTURE RESEARCH

This paper is conceptual and practice-based. The four cases serve as analytical vignettes rather than evidence of causal effectiveness, and no claim is made that the proposed designs produced measurable gains. The Target–Ecology–Gap–AI Affordance–Constraint–Agency sequence should therefore be treated as a pedagogical heuristic requiring empirical examination.

Future research could investigate the framework through longitudinal multiple-case designs, learner–AI interaction logs, written artefacts, stimulated-recall interviews, discourse analysis, and measures of subsequent independent performance.

A central question is whether stronger language produced during AI interaction transfers when support is removed. Research should also examine how productive constraints should be calibrated and sequenced as competence develops; whether AI-mediated target environments function differently across Modern Standard Arabic, academic and professional registers, religious discourse, and regional spoken varieties; and how teacher regulation develops into learners' capacity to design their own AI-mediated environments.

The framework should also be tested beyond Arabic in multilingual higher-education settings where learners have substantial exposure to a language but uneven access to the forms of participation required by their academic or professional goals.

8. CONCLUSIONS

Generative AI changes Arabic language education not merely by adding tools, but by altering who can perform linguistic activity, when, and under what conditions. The four cases show why its pedagogical value cannot be inferred from the amount of AI used.

For an advanced educator outside an Arabic-speaking environment, GenAI can reduce the organisational barriers to sustained meaningful use. For an

intermediate graduate student, it can provide recurrent, low-risk everyday production with limited scaffolding. For an Al-Azhar graduate student, it can add a productive academic layer that years of geographical immersion have not supplied. For a professional translator, by contrast, some of its most powerful generative functions must be restricted because independent writing is the competence under development.

The central question is not “How much should learners use AI?” but “What should learners delegate to AI in relation to what they are trying to learn to do themselves?” This shifts attention from AI use to the allocation of linguistic activity and redefines the teacher’s role.

The teacher increasingly designs relationships among learner needs, target activities, existing language ecologies, technological affordances, productive constraints, and agency. Such design is particularly important in Arabic, where geographical immersion, exposure to Standard Arabic, dialect use, and access to specialised registers may diverge sharply.

The goal is neither to minimise AI use nor to preserve effort for its own sake. It is to enlarge learners’ opportunities for meaningful linguistic participation without allowing the technology to appropriate the competence they are trying to build.

REFERENCES

- Albirini, A., et al. (2025). A systematic review of teaching and learning in the context of Arabic diglossia. *Social Sciences & Humanities Open*, 12, 102281. <https://doi.org/10.1016/j.ssaho.2025.102281>
- Generative AI Implementation and Assessment in Arabic Language Teaching. (2025). *International Journal of Online Pedagogy and Course Design*. <https://doi.org/10.4018/IJOPCD.368037>
- Javadinejad, A., & Davari, M. (2026). From tools to pedagogical ecologies: A systematic review of artificial intelligence in EFL/ESL education. *Frontiers in Education*, 11. <https://doi.org/10.3389/feduc.2026.1859248>
- Jia, W., & Pun, J. (2025). Exploring EMI teachers’ agency in addressing language-related challenges with ChatGPT: A multiple case study. *Computer Assisted Language Learning*. <https://doi.org/10.1080/09588221.2025.2511061>
- Lai, C., Pan, L., Pi, Z., Shen, T., & Liu, M. (2026). Learner agency in self-directed use of GenAI for informal language learning. *Innovation in Language Learning and Teaching*. <https://doi.org/10.1080/17501229.2026.2648781>
- Liu, Q., & Chao, C.-C. (2018). CALL from an ecological perspective: How a teacher perceives affordance and fosters learner agency in a technology-mediated language classroom. *ReCALL*, 30(1), 68–87. <https://doi.org/10.1017/S0958344017000222>

Long, M. H. (Ed.). (2005). *Second Language Needs Analysis*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511667299>

Morton, E., & Peeters, W. (2026). Orchestrating ChatGPT in English language teaching: A complex dynamic systems study of teacher agency. *Innovation in Language Learning and Teaching*. <https://doi.org/10.1080/17501229.2026.2708158>

Salsabila, A. D., Fikroh, A. N., & Bazighoh, Z. Z. (2026). Pembelajaran Bahasa Arab Untuk Tujuan Khusus (Arabic for Specific Purposes): Tinjauan sistematis dalam konteks pendidikan. *Mahira: Journal of Arabic Studies & Teaching*, 4(1). <https://doi.org/10.14421/mahira.2026.41.01>

Wang, G., Wang, W., Yang, D., & Ren, J. (2026). Generative AI, cognitive offloading, and learner agency in higher education: A scoping review. *Behavioral Sciences*, 16(7), 1150. <https://doi.org/10.3390/bs16071150>