Bridging Generative AI and Gamification in Moodle: Design and Evaluation of GamiBot

José Carlos Paiva
Pedagogical Innovation Center, Polytechnic of Porto, Portugal

Ricardo Queirós
CRACS, INESC TEC & ESMAD, Polytechnic of Porto, Portugal

Raquel Simões de Almeida
CIR, ESS, Polytechnic of Porto, Portugal

Teresa Terroso
ID+, ESMAD, Polytechnic of Porto, Portugal

Mário Pinto
ID+, ESMAD, Polytechnic of Porto, Portugal

Abstract

While Learning Management Systems (LMS) like Moodle are ubiquitous in higher education, they often lack dynamic, personalized academic support. We present GamiBot, an intelligent pedagogical agent natively integrated into Moodle to bridge this gap. GamiBot combines Large Language Model (LLM) assistance, Retrieval-Augmented Generation (RAG) grounded in course materials, adaptive quizzes, and a structured gamification backend. To evaluate its effectiveness, we conducted a two-week pilot study across two higher education institutions. Results from telemetry (773 interactions) and student surveys demonstrated high system usability (74.6% positive sentiment) and strong chatbot performance (69.4%). Students heavily favored active self-regulation, dedicating 53% of interactions to generating adaptive quizzes rather than passive content summaries. Conversely, gamification elements received mixed feedback (31.7%), underscoring the need for technical refinement and high customizability. Ultimately, GamiBot suggests that integrating generative AI directly into existing LMS workflows can provide reliable, context-aware formative support that aligns with self-regulated learning strategies.

2012 ACM Subject Classification Applied computing → Education; Applied computing → E-learning; Human-centered computing → Interactive systems and tools; Computing methodologies → Artificial intelligence; Human-centered computing → HCI design and evaluation methods; Human-centered computing → Collaborative and social computing systems and tools

Keywords and phrases Generative AI, Educational chatbot, Pedagogical agent, Gamification, Moodle, Retrieval-augmented generation, Self-regulated learning

Digital Object Identifier 10.4230/OASIcs.ICPEC.2026.17

Supplementary Material
Text (Project Website): `https://inov-norte.github.io/gamibot-docs` [18]

Funding *Ricardo Queirós*: This work is funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under the support UID/50014/2025 (`https://doi.org/10.54499/UID/50014/2025`).
Raquel Simões de Almeida: This work has been supported by Fundação para a Ciência e Tecnologia (FCT) through R&D Units funding (UID/5210/2025) `https://doi.org/10.54499/UID/05210/2025`.
Teresa Terroso: Funded by FCT – Fundação para a Ciência e a Tecnologia (UID/04057/2025, `https://doi.org/10.54499/UID/04057/2025`).

1 Introduction

Student engagement and self-regulated learning represent persistent challenges in higher education, particularly within technology-mediated environments. Learning Management Systems (LMS) such as Moodle have become cornerstones of institutional education, providing a centralized space for course materials, assessments, and communication [8]. Yet, despite their widespread adoption, these platforms often fall short in offering dynamic, personalized academic support. Students navigating complex course content must typically rely on static resources, delayed instructor feedback, or external tools that operate in isolation from their learning environment. This fragmentation between the learning environment and support tools creates friction that can undermine both motivation and academic performance [16].

In parallel, two prominent approaches have emerged to address engagement and support in digital learning: gamification and artificial intelligence (AI)-powered pedagogical agents. Gamification – the application of game design elements such as points, badges, and leaderboards to non-game contexts – has demonstrated potential for increasing student motivation and participation [10]. Meanwhile, AI chatbots and conversational agents have shown promise in providing on-demand tutoring, content clarification, and formative assessment at scale [4]. However, these two paradigms are rarely combined in a meaningful, integrated manner, and even more rarely deployed as native components of institutional LMS platforms. Most existing solutions remain either standalone applications or loosely coupled external tools, leaving a critical gap between the capabilities of modern generative AI and the institutional workflows in which students and teachers actually operate.

This paper addresses that gap by presenting GamiBot, an intelligent, gamified pedagogical chatbot designed for deep, native integration within Moodle. GamiBot combines large language model (LLM)-based assistance with retrieval-augmented generation (RAG) over course-specific materials, adaptive quiz generation, and a structured gamification layer – all delivered through a unified conversational interface embedded directly within the LMS. Rather than requiring students to switch between platforms or teachers to maintain separate tools, GamiBot situates AI-powered support within the institutional context where learning already occurs, keeping interactions anchored to official course content and pedagogical workflows. The project documentation and public website are available at `https://inov-norte.github.io/gamibot-docs`.

The system was architected around three interdependent objectives: delivering contextually grounded academic support through intelligent material summarization and instant answers to questions; enabling formative self-assessment through dynamically adaptive quizzes; and sustaining learner engagement through a coherent, progression-driven gamification framework. To evaluate whether these objectives translate into real pedagogical value, we conducted an empirical study involving both students and teachers. Specifically, this paper seeks to answer the following research questions:

- **RQ1**: How do students perceive GamiBot in terms of usefulness, engagement, and ease of use?
- **RQ2**: How do teachers perceive its pedagogical value and classroom applicability?
- **RQ3**: How do students' interaction patterns (telemetry) reflect self-regulated learning behaviors?

The remainder of this paper is organized as follows. Section 2 reviews related work on gamification in education, AI pedagogical agents, and the LMS integration gap. Section 3 details the design and implementation of GamiBot, covering its architecture, pedagogical

features, gamification mechanics, and user interface. Section 4 presents the evaluation methodology and results. Section 5 discusses the findings in relation to the research questions, and Section 6 concludes with implications and directions for future work.

2 Related Work

The development of GamiBot draws on four interconnected bodies of literature: the application of gamification in educational contexts, the design and evaluation of AI-powered chatbots and pedagogical agents, emerging work on systems that combine both paradigms, and the persistent challenge of integrating these solutions into institutional learning management systems.

2.1 Gamification in Education

Gamification – broadly defined as the use of game design elements in non-game contexts [5] – has been extensively studied as a mechanism for enhancing student motivation, engagement, and persistence in educational settings. Early empirical work demonstrated that the introduction of points, badges, and leaderboards into classroom activities could positively influence participation and task completion rates [9]. Subsequent research, however, introduced important nuances: the effectiveness of gamification is neither universal nor unconditional, and depends significantly on the meaningfulness of the reward structures, the alignment between game mechanics and learning objectives, and individual differences in learner profiles such as achievement orientation and intrinsic motivation [15, 22]. Studies have consistently warned against superficial implementations – so-called "pointsification" – in which rewards are decoupled from substantive academic actions, resulting in short-term compliance rather than deep engagement [17]. More recent work has advocated for adaptive and personalized gamification designs that respond to individual learner behavior and progression over time [12], a principle that informs the tiered badge system and event-driven reward mechanics adopted in GamiBot.

2.2 AI Chatbots and Pedagogical Agents

The use of conversational agents in educational settings has a long history, originating in intelligent tutoring systems (ITS) that modeled learner knowledge states and adapted instructional sequences accordingly [1, 21]. The advent of large language models has substantially expanded the capabilities and accessibility of such systems. Contemporary LLM-based educational chatbots are capable of handling open-ended natural language queries, generating explanations at varying levels of abstraction, and producing formative assessment items on demand, without the need for laboriously hand-crafted knowledge bases [11]. Retrieval-augmented generation (RAG) architectures have further strengthened these systems by grounding model outputs in domain-specific or course-specific document collections, reducing hallucination and increasing the contextual relevance of responses [14]. Empirical evaluations of LLM-based pedagogical agents have reported positive effects on students' perceived usefulness, learning efficiency, and willingness to engage with course content [11, 2], though concerns regarding accuracy, over-reliance, and the displacement of critical thinking have also been raised [7]. GamiBot builds on this tradition by combining RAG-grounded assistance with adaptive quiz generation, ensuring that AI-generated content remains tethered to institutionally curated materials rather than the unconstrained parametric knowledge of the underlying model.

2.3 Gamified Educational Agents Combining AI and Gamification

Despite the independent maturity of both gamification research and AI-driven pedagogical agents, systems that meaningfully integrate both paradigms remain relatively scarce in the literature. Several studies have explored the sequential or additive combination of chatbots with gamified elements – for instance, embedding chatbot interactions within gamified course shells or awarding points for chatbot usage – but these approaches typically treat the two components as independent modules rather than as a coherent, mutually reinforcing design [19]. A smaller set of works has proposed more tightly coupled architectures in which conversational agents adapt their behavior based on learner progression data, or in which gamification rewards are triggered dynamically by the nature and quality of the agent interaction [13]. The evidence base for such integrated systems remains limited, and existing prototypes have rarely been evaluated empirically in authentic classroom settings. GamiBot contributes to this emerging line of work by implementing a dedicated gamification backend that responds to semantically meaningful pedagogical actions, including question asking, summarization requests, and quiz engagement, thereby linking motivational incentives directly to the learning behaviors the system is designed to encourage.

2.4 The LMS Integration Gap

A recurring limitation in the educational technology literature concerns the difficulty of embedding AI and gamification tools within the institutional platforms that teachers and students already use on a daily basis. Learning management systems such as Moodle, Canvas, and Blackboard serve as the primary digital infrastructure for course delivery in higher education, yet they offer limited native support for generative AI or sophisticated gamification mechanics [20]. As a consequence, most AI-powered educational tools are deployed as standalone web applications or third-party plugins that require students to leave the LMS environment to access support, disrupting the continuity of the learning experience and reducing the likelihood of adoption [23]. Similarly, gamification in LMS contexts has often been implemented through external platforms such as Kahoot or Classcraft, which operate independently of the course content and activity data available within the LMS [6]. This fragmentation has practical consequences: teachers must maintain parallel systems, students must context-switch between tools, and the behavioral and performance data generated by AI or gamification interactions remains disconnected from the institutional record. GamiBot is specifically designed to close this integration gap by functioning as a native Moodle extension, co-locating AI assistance, adaptive assessment, and gamification feedback within the same interface through which students already access their course materials and activities.

3 GamiBot

To address the documented gap between generative AI capabilities and institutional learning management systems, we developed GamiBot: an intelligent, gamified pedagogical agent designed for deep integration within Moodle. This section details the technical and pedagogical foundations of the proposed system. First, subsection 3.1 outlines the high-level system architecture, explaining the orchestration of its core components and the flow of data during user interactions. Next, subsection 3.2 describes the AI-driven pedagogical features, specifically focusing on real-time assistance and adaptive quiz generation. Subsection 3.3 details the gamification design and its underlying mechanics. Finally, subsection 3.4 presents the user interface, explaining how progression mechanics are embedded directly into the conversational workflow to foster engagement and self-regulated learning.

3.1 System Overview

GamiBot is a modular agent-based ecosystem designed to extend Moodle with context-aware generative AI support and embedded gamification mechanisms. In contrast with standalone educational chatbots, the system is conceived as an LMS-native solution, enabling students to access tutoring, summarization, and formative assessment directly inside the course environment rather than through external platforms. This design choice is pedagogically relevant because it keeps support interactions close to the learning materials, the ongoing activities, and the institutional workflows already used by teachers and students.

At the architectural level, as illustrated in Figure 1, GamiBot is organized into four interconnected layers:

Client Layer: composed of the Moodle LMS and the GamiBot plugins, this layer provides the unified chat and gamification interface. It serves as the primary entry point for both students and professors, mediating all platform integrations.

Orchestration Layer: implemented via LangFlow, this layer is responsible for the coordination of pedagogical agents. It receives incoming requests (via webhooks and APIs) and orchestrates the workflows required to formulate a response.

Intelligence Layer: this layer executes the core cognitive tasks. It includes the Large Language Model (LLM) for text generation and the Qdrant vector database, which handles semantic search over course resources.

Gamification Layer: powered by GamiAPI, this layer manages learner progression, receiving user actions from the Client Layer to independently calculate and store rewards, points, and badges.

This separation of concerns allows the system to combine pedagogical flexibility with technical maintainability. Moodle remains the primary learning environment, LangFlow coordinates intelligent behavior, Qdrant and the LLM support content generation, and GamiAPI handles gamification logic independently of the conversational flow.

The Intelligence Layer is powered by OpenAI's `gpt-oss-20b`, a 20-billion parameter open-weights large language model. The selection of an openly available model is particularly advantageous in an educational context, as it allows us (and other institutions) to self-host the infrastructure. This ensures that sensitive student interactions and proprietary course materials remain secure and compliant with institutional data privacy regulations. This model serves as the underlying reasoning engine for both the conversational assistant and the adaptive quiz-generation workflow, operating over prompts enriched with retrieved course context.

A key feature of GamiBot architecture is the continuous alignment between course content and AI support. As instructors build or update a Moodle course, educational materials such as PDFs, slide decks, textual resources, and other uploaded assets are processed and indexed in Qdrant through vector embeddings. As a result, agents within the Orchestration Layer can later retrieve semantically relevant fragments of official course materials without requiring teachers to manually author question–answer pairs or scripted tutoring rules. This mechanism ensures that the system remains synchronized with the evolving course and that its outputs are anchored in institutionally curated content rather than relying exclusively on the parametric knowledge of the underlying LLM.

Operationally, when a student submits a question or requests a quiz through the interface, the request is passed by the Client Layer to LangFlow together with relevant contextual metadata. The active agent in the Orchestration Layer then determines the appropriate pedagogical action, queries Qdrant for semantic search when necessary, composes an enriched

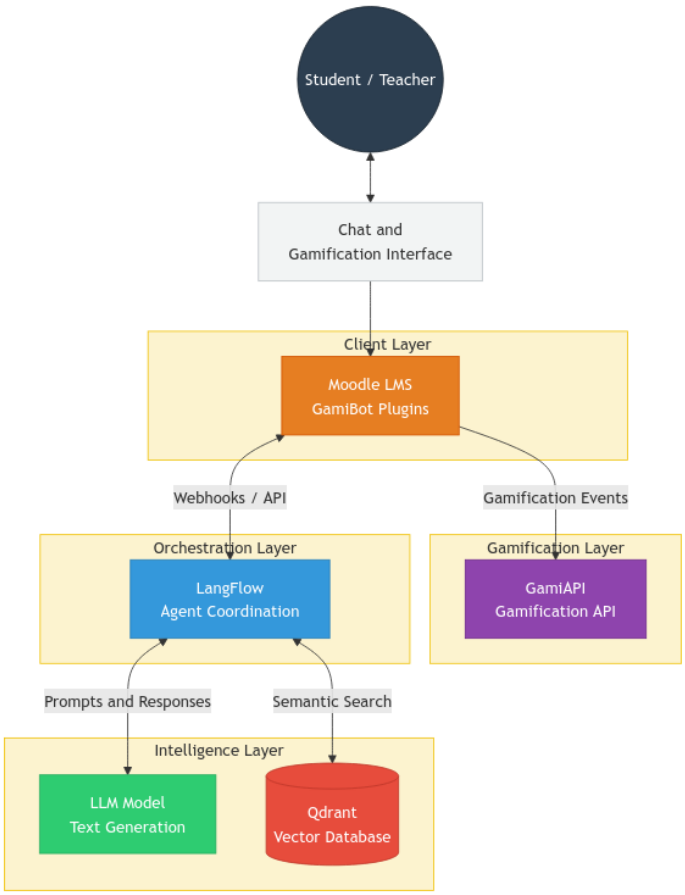

Figure 1 High-level architecture of the GamiBot ecosystem, illustrating the data flow and interactions across the Client, Orchestration, Intelligence, and Gamification layers.

prompt, and forwards it to the LLM in the Intelligence Layer. In parallel, the Client Layer reports the corresponding user action to GamiAPI in the Gamification Layer, which updates the learner's state and returns the refreshed profile to Moodle. The student therefore receives both the AI-generated educational output and the updated progress indicators within a single integrated interface.

3.2 Pedagogical Features

GamiBot provides context-aware support through two core AI-driven features: real-time assistance and adaptive quiz generation. Central to both workflows is the `gpt-oss-20b` large language model. By injecting explicit pedagogical instructions alongside semantically relevant course excerpts retrieved from Qdrant into the model's context window, these features position the system not merely as a conversational interface, but as a pedagogical assistant capable of responding to immediate learning needs while encouraging active engagement.

The first feature, real-time assistance and summarization, supports students during their interaction with Moodle resources. Learners may ask questions about a topic, request clarification about a specific concept, or ask for the summary of a section they are studying. In these cases, an assistance agent retrieves relevant content from Qdrant and uses it to

construct a grounded prompt for the `gpt-oss-20b` model. The resulting response is strictly anchored to the institutional curriculum, which increases contextual relevance and reduces the risk of generic or decontextualized explanations.

The second feature, adaptive quiz generation, supports formative assessment and knowledge consolidation. Rather than relying on a fixed bank of manually authored questions, GamiBot uses a dedicated quiz-generation agent that invokes Qdrant during its reasoning process. When a quiz is requested, the agent retrieves topic-specific content segments. The system prompt then constrains the `gpt-oss-20b` model to generate formative assessment items strictly aligned with this retrieved evidence. This dynamic generation minimizes hallucinations and allows the quizzes to remain sensitive to the course context while significantly reducing authoring overhead for instructors.

An important aspect of this feature is that the quiz interaction does not end with answer submission. After the student responds, the system produces a structured performance summary that includes the submitted answers, the expected correct answers, explanatory feedback, and targeted suggestions for further study. In pedagogical terms, this transforms quiz generation from a simple testing mechanism into a feedback-oriented formative loop. The value of the interaction therefore lies not only in whether the student answered correctly, but also in the interpretive layer that helps them understand mistakes and identify areas requiring further attention.

Combined, these two features support a cycle of inquiry, feedback, and reinforcement. Students can begin with a question, receive grounded assistance, test their understanding through an adaptive quiz, and then review a personalized explanation of their performance. This sequence is particularly compatible with self-regulated learning processes, as it encourages learners to identify knowledge gaps, seek targeted support, and act on feedback within the same digital environment.

3.3 Gamification Design

Beyond its AI capabilities, GamiBot incorporates a dedicated gamification layer intended to transform learning interactions into visible and motivating progress. Rather than treating gamification as a static overlay, the system links rewards to meaningful pedagogical actions performed within the conversational workflow. Through interaction with GamiBot, students accumulate points for actions such as asking questions, requesting summaries, or engaging with quizzes; earn badges associated with distinct forms of participation; and progress through topics, reflecting their continued engagement over the duration of the course.

To provide granular and continuous recognition, badges in GamiBot are structured around a three-level progression system. This multi-tier approach ensures that students remain motivated even after unlocking an initial reward, preventing the sudden drop-off in engagement often associated with one-time achievements. Examples of this badge progression include:

- The **Educational Material Analyst** badge, which recognizes active doubt clarification. A student unlocks Level 1 after asking 3 context-related questions, upgrades to Level 2 at 10 questions, and reaches Level 3 mastery at 20 questions.
- The **Rough Diamond** badge, which rewards persistence and the willingness to learn from mistakes. Level 1 is awarded when a student passes multiple quizzes after initially submitting incorrect answers, advancing to Levels 2 and 3 as they continually polish their understanding through trial, error, and reflection.
- The **Quiz Master** badge, which acknowledges consistent engagement with formative assessment. The first tier is awarded after completing 3 quizzes, scaling up to higher tiers as the learner routinely tests their knowledge across different course topics.

Furthermore, the gamification design includes a quiz life mechanic to reinforce purposeful engagement. Each quiz session begins with a limited number of lives (typically three), and incorrect answers consume one life. By making attempts finite, this mechanic discourages random guessing and encourages students to reflect carefully before responding or to revisit the AI's explanatory support when uncertain.

All of these mechanics are managed by GamiAPI, which operates as a separate backend service responsible for tracking user state, validating events, and returning updated progression data to Moodle. This architectural design has two distinct advantages. First, it keeps gamification rules configurable and reusable across courses or deployments without tightly coupling them to the conversational logic. Second, it allows the platform to record richer activity traces, including point accumulation rates, badge tier upgrades, and life consumption metrics, which can later support both learner self-monitoring and empirical analysis of student behavior.

3.4 User Interface

The user interface is intentionally chat-centric, designed to deliver motivational feedback synchronously with AI-driven academic support. From the student's perspective, GamiBot appears as an integrated conversational component inside Moodle, combining AI responses, interactive assessments, and gamification feedback within the same interaction space. This co-location is important as it ensures that pedagogical help and motivational feedback are not experienced as separate systems, but as logical parts of a single continuous learning interaction.

As illustrated in Figure 2, the interface is divided into distinct contextual views that guide the learner through the pedagogical loop:

- **Welcome and Action Menu**: upon opening GamiBot, the student is greeted with quick-action buttons to generate quizzes, request summaries, clarify doubts, or view their progress. This minimizes the cognitive load required to initiate a learning session.
- **Conversational Assistance**: the chat window allows for natural language queries. The AI's responses explicitly reference the course materials (e.g., citing specific documents) to ensure the student knows the information is grounded in the official curriculum.
- **Adaptive Quizzes and Feedback**: when a student engages in a quiz, the chat transforms into a testing interface. The header displays the remaining quiz lives (represented as hearts), enforcing the finite attempt mechanic. Immediate feedback is provided after each question, offering detailed justifications for both correct and incorrect answers. At the end of the quiz, the system automatically generates a personalized study guide highlighting the concepts the student struggled with, areas for improvement, and tips.
- **Progress Dashboard**: the progress area exposes detailed usage statistics (e.g., points, earned badges, time spent, and number of quizzes completed). It visually displays the tiered badges the student has unlocked (such as the "Rough Diamond" badge) and highlights specific topics where the student has struggled, enabling them to monitor their engagement and identify areas where further study may be needed.

Overall, GamiBot combines Moodle integration, agent orchestration, retrieval-augmented generation, and gamification into a cohesive educational ecosystem. Its contribution lies not only in adding AI features to an LMS, but in structuring these features around grounded support, adaptive formative assessment, and visible progression, thereby creating a more interactive and pedagogically expressive digital learning environment.

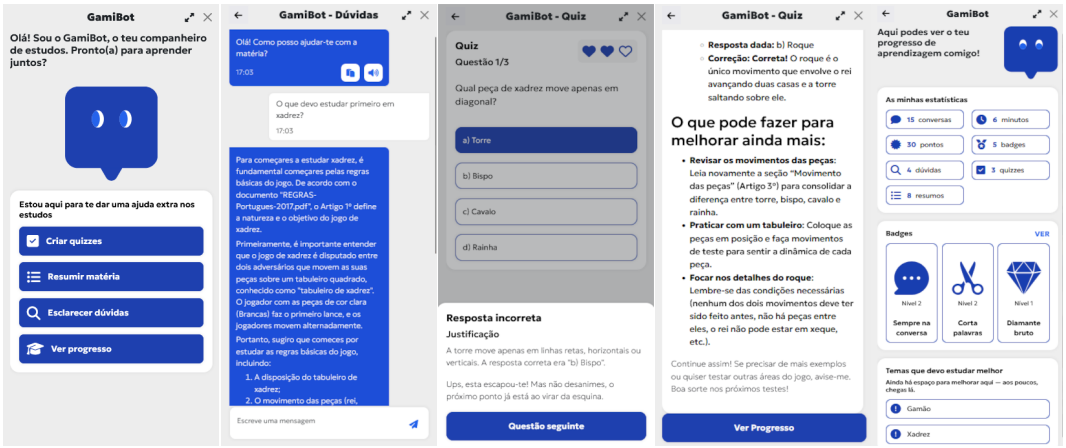

Figure 2 Screenshots of the GamiBot interface illustrating its core functionalities: quick-action menus, retrieval-augmented doubt clarification, adaptive quizzes with a finite-life mechanic, immediate formative feedback, and a comprehensive gamification dashboard that tracks student progress and achievements.

4 Evaluation

To assess the usability, pedagogical value, and engagement impact of GamiBot in real educational contexts, we conducted an exploratory pilot study across two higher education institutions within the Polytechnic Institute of Porto (IPP). This section details the methodology used to collect data and presents the empirical results of the intervention.

4.1 Context and Methodology

GamiBot was deployed directly into the Moodle environments of the School of Health (ESS IPP) and the School of Media Arts and Design (ESMAD IPP). To test the system's versatility across different knowledge domains and academic levels, it was activated in four distinct curricular units:

- **ESS IPP:** Master's in Occupational Therapy (*Neuroscience of Mental Illness*) and Master's in Digital Transformation of Health and Rehabilitation Services (*Serious Games and Gamification in Health and Rehabilitation*).
- **ESMAD IPP:** Bachelor's in Web Information Systems and Technologies (*Object-Oriented Programming*, 1st year; and *Web Programming II*, 2nd year).

The target population comprised 164 students (127 from ESMAD and 37 from ESS), resulting in a self-selected sample of 32 participants (19.5% response rate) for the survey evaluation. Participation was entirely voluntary, with the system positioned as a supplementary study tool rather than a mandatory course requirement.

Data were collected using a mixed-methods approach over a two-week pilot period. First, we extracted quantitative telemetry from the GamiBot orchestration layer to measure actual usage, tracking the volume of AI-generated quizzes, clarified doubts, and generated summaries. Second, at the end of the pilot period, students were invited to complete an anonymous, non-mandatory online questionnaire. The survey utilized a 5-point Likert scale (ranging from "Strongly Disagree" to "Strongly Agree") to evaluate five primary dimensions:

- **Usability:** assessed using items adapted from the System Usability Scale (SUS).
- **Chatbot Experience:** evaluating the clarity, contextual relevance, and accuracy of the AI's natural language responses.
- **Gamification:** measuring the motivational impact of points, badges, and progression levels.
- **Pedagogical Impact:** assessing perceived improvements in content comprehension and learner autonomy.
- **Intention to Use:** based on constructs from the Technology Acceptance Model (TAM), measuring the likelihood of future adoption and peer recommendation.

Open-ended questions were also included at the end of the survey to capture qualitative feedback on what students liked most about the system and what specific areas required improvement.

During this pilot, all AI-generated quizzes, doubt clarifications, and summaries were powered by the `gpt-oss-20b` model. Due to initial local hardware constraints, the model was deployed as a managed online endpoint via Microsoft Azure AI Foundry for the duration of the experiment. However, because `gpt-oss-20b` is an open-weights model, the architecture remains fully compatible with on-premise self-hosting for future, permanent institutional deployments.

4.2 System Usage and Interaction Patterns

Log data from the pilot revealed active, voluntary engagement from the student body across both institutions. Overall, the system processed a total of 773 distinct pedagogical interactions.

At ESMAD, students generated 361 quizzes, requested 142 course summaries, and asked 49 specific questions (doubts). At ESS, despite the smaller cohort, engagement remained high relative to class size, with 48 generated quizzes, 45 summaries, and 28 clarified doubts. Notably, across both institutions, the generation of adaptive quizzes was by far the most heavily utilized feature, accounting for approximately 53% of all interactions. This suggests a strong student preference for active, formative assessment over passive content summarization when using an AI assistant.

4.3 Student Perspectives on Usability and Pedagogy

A total of 32 students responded to the voluntary evaluation questionnaire. Of these, the vast majority (81%) reported using GamiBot at least once, with nearly half (47%) using it multiple times during the short pilot.

The student evaluation across the five primary dimensions yielded very encouraging results, particularly regarding the core system functionalities and the pedagogical value of the agent. Figure 3 summarizes the breakdown of Negative, Neutral, and Positive responses across all 29 survey items, properly accounting for reverse-coded statements in the usability scale.

Students reported exceptionally high levels of usability. Across the ten items adapted from the System Usability Scale (SUS), the overall positive sentiment was 74.6%. When asked specifically if the system was easy to use, 83.4% agreed or strongly agreed, highlighting the success of integrating the tool directly into the Moodle interface.

In terms of the Chatbot Experience, 69.4% of the comprehensive responses were positive. Specifically, 73.4% agreed or strongly agreed that GamiBot's answers were clear and understandable, and 70% agreed that the system responded adequately to the specific context of their discipline, validating the effectiveness of the Qdrant-based retrieval-augmented generation (RAG) approach.

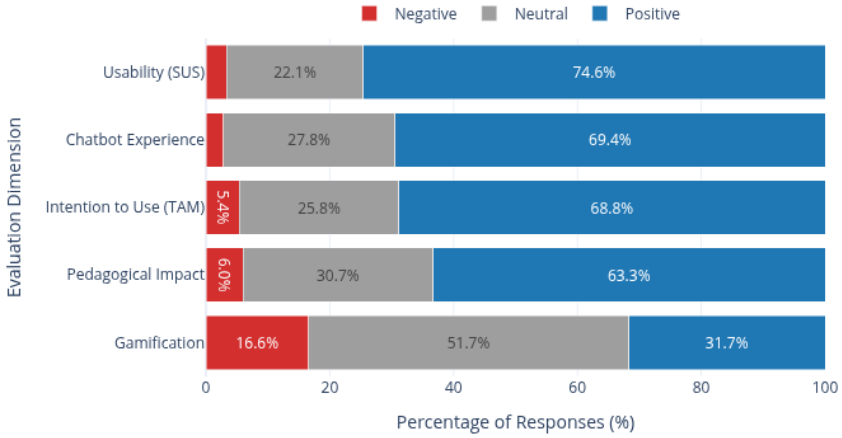

Figure 3 Student evaluation of GamiBot across five primary dimensions, showing the percentage of Negative (Disagree/Strongly Disagree), Neutral, and Positive (Agree/Strongly Agree) responses (n=32).

Regarding Pedagogical Impact, the overall positive sentiment reached 63.3%. Individually, 66.7% of respondents considered GamiBot a useful tool within the context of their curricular unit, and 56.7% agreed or strongly agreed that it improved their autonomy in learning. This sentiment was echoed in the qualitative open-ended feedback, where one student noted: *"[It is] highly intuitive, extremely easy to use; the responses are precise and well contextualized with the material provided in Moodle."*

Furthermore, the Intention to Use dimension achieved a 68.8% positive evaluation. Looking forward, 70% of the respondents stated they intend to continue using GamiBot, and nearly 60% said they would recommend the tool to their colleagues.

4.4 Perspectives on Gamification Mechanics

While the AI and usability aspects were highly rated, feedback on the gamification elements was more mixed. Across the five items measuring this dimension, the overall positive sentiment was 31.7%. When asked specifically if the experience points motivated them to continue using the system, 37.9% agreed or strongly agreed, but a significant portion (48.3%) remained neutral ("Neither agree nor disagree"), and 13.7% disagreed. Similar distributions were observed regarding the impact of badges on making the learning experience more interesting.

Qualitative feedback provided crucial context for these mixed results. Several students reported technical glitches with the progression interface during the pilot, stating, *"The gamification functionality was not working properly [at times], which is why I answered 'strongly disagree'."*. Additionally, students expressed a desire for greater control over the gamified elements, such as being able to explicitly choose the number of questions in a quiz session rather than relying on a fixed length set by teachers.

Despite these interface bugs, the overall reception of the core tool was highly positive, establishing a strong foundation for future iterative improvements to the gamified components.

5 Discussion

Evaluating GamiBot across multiple academic disciplines provided valuable insights into the integration of AI-driven, gamified pedagogical agents within LMSs. Addressing our first research question (RQ1) regarding student perceptions, the system was highly successful

in terms of ease of use and usefulness. The seamless integration directly into the Moodle interface eliminated the friction of navigating to third-party platforms, which was reflected in the 74.6% positive sentiment in the Usability (SUS) dimension, with 83.4% of students explicitly agreeing that the system was easy to use. In terms of usefulness, the Chatbot Experience dimension received 69.4% positive feedback. Students praised the clarity and contextual relevance of the AI's responses, confirming that the RAG architecture successfully grounded the language model in course-specific materials. This reliable performance built student trust and drove a high Intention to Use (68.8% positive).

However, student perception of engagement through gamification yielded mixed results (31.7% positive sentiment). While the core AI features effectively engaged students, the specific gamified elements did not achieve their intended motivational impact for the majority. Qualitative feedback revealed this was largely due to technical glitches in the progression interface and a lack of user control over gamified tasks, such as fixed quiz lengths. This dichotomy suggests that while students highly value reliable AI assistance, gamification in higher education must transcend "pointsification" and ensure high technical stability to avoid becoming a source of cognitive friction rather than an engaging motivator.

Regarding the second research question (RQ2) on pedagogical value and classroom applicability, the system demonstrated high versatility by functioning effectively across different knowledge domains. It adapted seamlessly to both technical, hard-skill courses and theoretical, health-science courses. Because the RAG pipeline utilizes existing Moodle documents, teachers were not required to manually train the AI, drastically lowering the barrier to classroom adoption. Furthermore, by autonomously processing 773 pedagogical interactions over the short pilot, including answering 77 specific student doubts and generating 409 targeted quizzes, GamiBot effectively acted as a 24/7 supplementary teaching assistant. This capacity to handle high-volume, low-complexity queries suggests strong pedagogical value, allowing educators to dedicate more classroom time to complex problem-solving and critical discussion.

Addressing the third research question (RQ3), the empirical data provides compelling evidence of GamiBot's positive impact on participation, self-regulation, and learning support. Participation was strong, with 81% of surveyed students interacting with the bot voluntarily and 47% returning to use it multiple times. Crucially, the telemetry revealed a distinct trend toward self-regulated learning: approximately 53% of all system interactions were requests for adaptive quizzes, far outpacing passive requests such as course summaries or direct Q&A. This behavioral pattern is consistent with the formative assessment framework proposed by Black and Wiliam [3], particularly the strategy of activating students as owners of their own learning. By prioritizing generative testing over passive review, learners used the agent to close the gap between their current understanding and the intended learning goals, thereby transforming the LMS into a responsive environment for self-regulated growth. This interpretation is further supported by the survey results on learning support, where 63.3% of responses in the Pedagogical Impact dimension were positive, and 56.7% of students explicitly agreed that the system improved their learning autonomy.

While these initial results are promising, this pilot study has limitations. The sample size for the evaluation questionnaire (n=32) was relatively small, and the pilot duration was brief. Additionally, while the usage data strongly implies benefits for educators, future studies must include formal qualitative and quantitative evaluations directly from the teachers to fully validate their perceptions of the system's classroom applicability.

6 Conclusion and Future Work

This study presented GamiBot, an educational agent that bridges the gap between generative AI, gamification, and institutional LMSs. The pilot deployment demonstrated that embedding a RAG-powered LLM directly within Moodle provides highly accessible and contextually accurate academic support without forcing users to leave their learning environment. Telemetry and survey data revealed that students heavily favor active self-assessment, frequently utilizing the system to generate adaptive quizzes. This underscores GamiBot's significant pedagogical value in promoting learner autonomy and self-regulated learning.

Conversely, the mixed reception of the gamification elements highlighted a critical design lesson: to be effective, motivational mechanics must be technically robust and highly adaptable to individual preferences. Future work will focus on stabilizing the gamification engine and granting students greater control over their learning parameters, such as dynamically adjusting quiz lengths and difficulty levels. Additionally, we plan to conduct longitudinal studies with larger cohorts to measure GamiBot's impact on objective academic performance, alongside formal evaluations to capture the long-term benefits for educators managing the platform.

References

1 John R Anderson, C Franklin Boyle, and Brian J Reiser. Intelligent tutoring systems. *Science*, 228(4698):456–462, 1985.

2 David Baidoo-Anu and Leticia Owusu Ansah. Education in the era of generative artificial intelligence (ChatGPT): Understanding the potential benefits and challenges. *Journal of AI*, 7(1):52–62, 2023.

3 Paul Black and Dylan Wiliam. Developing the theory of formative assessment. *Educational Assessment, Evaluation and Accountability(formerly: Journal of Personnel Evaluation in Education)*, 21(1):5–31, February 2009. doi:10.1007/s11092-008-9068-5.

4 Narius Farhad Davar, M. Ali Akber Dewan, and Xiaokun Zhang. Ai chatbots in education: Challenges and opportunities. *Information*, 16(3), 2025. doi:10.3390/info16030235.

5 Sebastian Deterding, Dan Dixon, Rilla Khaled, and Lennart Nacke. From game design elements to gamefulness: defining gamification. In *Proceedings of the 15th International Academic MindTrek Conference*, pages 9–15. ACM, 2011. doi:10.1145/2181037.2181040.

6 Darina Dicheva, Christo Dichev, Gennady Agre, and Gabriela Angelova. Gamification in education: A systematic mapping study. *Journal of Educational Technology & Society*, 18(3):75–88, 2015. URL: https://www.j-ets.net/ETS/journals/18_3/6.pdf.

7 Yogesh K Dwivedi, Nir Kshetri, Laurie Hughes, Emma Louise Slade, Anand Jeyaraj, Arpan Kumar Kar, Abdullah M Baabdullah, Alex Koohang, Vishnupriya Raghavan, Manju Ahuja, et al. Opinion paper: "so what if chatgpt wrote it?" multidisciplinary perspectives on opportunities, challenges and implications of generative conversational ai for research, practice and policy. *International Journal of Information Management*, 71:102642, 2023. doi:10.1016/J.IJINFOMGT.2023.102642.

8 Sithara H. P. W. Gamage, Jennifer R. Ayres, and Monica B. Behrend. A systematic review on trends in using moodle for teaching and learning. *International Journal of STEM Education*, 9(1):9, January 2022. doi:10.1186/s40594-021-00323-x.

9 Juho Hamari, Jonna Koivisto, and Harri Sarsa. Does gamification work? a literature review of empirical studies on gamification. In *Proceedings of the 47th Hawaii International Conference on System Sciences (HICSS-47)*, pages 3025–3034. IEEE, 2014. doi:10.1109/HICSS.2014.377.

10 Carlos J. Hellín, Francisco Calles-Esteban, Adrián Valledor, Josefa Gómez, Salvador Otón-Tortosa, and Abdelhamid Tayebi. Enhancing student motivation and engagement through a gamified learning environment. *Sustainability*, 15(19), 2023. doi:10.3390/su151914119.

11 Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günnemann, Eyke Hüllermeier, et al. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103:102274, 2023.

12 Ana Carolina Tomé Klock, Isabela Gasparini, Marcelo Soares Pimenta, and Juho Hamari. Adaptive gamification in education: A literature review of adaptive frameworks, applications and instruments. *International Journal of Human-Computer Studies*, 141:102496, 2020.

13 Lela Labadze, Maya Grigolia, and Lasha Machaidze. Role of ai chatbots in education: Systematic literature review. *International Journal of Educational Technology in Higher Education*, 20(1):56, 2023.

14 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474, 2020.

15 Elisa D Mekler, Florian Brühlmann, Alexandre N Tuch, and Klaus Opwis. Towards understanding the effects of individual gamification elements on intrinsic motivation and performance. *Computers in Human Behavior*, 71:525–534, 2017. doi:10.1016/J.CHB.2015.08.048.

16 Negar Monazam-Tabrizi, Yusuf Kurt, and William il Kuk Kang. Navigating learning disruptions: The role of digital learning platforms in student motivation, feedback and emotion. *Computers & Education*, 246:105534, 2026. doi:10.1016/j.compedu.2025.105534.

17 Scott Nicholson. A recipe for meaningful gamification. In *Gamification in Education and Business*, pages 1–20. Springer, 2015.

18 José Paiva and Ricardo Queirós. Gamibot. Text (visited on 2026-07-13). URL: https://inov-norte.github.io/gamibot-docs, doi:10.4230/artifacts.27132.

19 Jorge Quiroga Pérez, Llúcia Daiño, Caterina Pinya, and Joan J Muntaner-Guasp. Chatbots and conversational agents in education: A bibliometric analysis using scopus database. *Interactive Learning Environments*, 31(7):4500–4515, 2020.

20 Ahmed Tlili, Boulus Shehata, Michael Agyemang Adarkwah, Aras Bozkurt, Daniel T Hickey, Ronghuai Huang, and Boulus Agyemang. If history repeats itself, how will ai reshape education? a literature review of the state of research on chatgpt in education. *Smart Learning Environments*, 10(1):32, 2023.

21 Kurt VanLehn. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4):197–221, 2011.

22 Maarten Vansteenkiste, Joke Simons, Willy Lens, Kennon M Sheldon, and Edward L Deci. Motivating learning, performance, and persistence: the synergistic effects of intrinsic goal content and autonomy-supportive context. *Journal of Personality and Social Psychology*, 87(2):246–260, 2004.

23 Olaf Zawacki-Richter, Victoria I Marín, Melissa Bond, and Franziska Gouverneur. Systematic review of research on artificial intelligence applications in higher education – where are the educators? *International Journal of Educational Technology in Higher Education*, 16(1):39, 2019.