Game-based and AI-assisted Learning about Quantum Science

Fabian Schrodt[1,1], Caroline Wermann[2,2], Francesco Catania[1,3],
Karina E. Avila[2,4], Maite Maurer[5,5], Sascha Mehlhase[5,6], Sebastian André[4,7],
Tatjana Wilk[4,8], Jochen Kuhn[2,9], Nikola Merkas[3,10], Stefan Küchemann[2,11]

[1]*Quantum Gaming GmbH, Tuebingen, Germany*
[2]*Chair of Physics Education Research, Ludwig-Maximilians-University of Munich, Germany*
[3]*Merkas Technologies GmbH, Zuerich, Switzerland*
[4]*Munich Center for Quantum Science and Technology, Munich, Germany*
[5]*Munich Quantum Valley, Munich, Germany*

[1]fabian@q-gaming.com, [2]caroline.wermann@lmu.de, [3]francesco@q-gaming.com,
[4]k.avila@physik.uni-muenchen.de, [5]maite.maurer@munich-quantum-valley.de,
[6]sascha.mehlhase@munich-quantum-valley.de, [7]s.andre@lmu.de, [8]tatjana.wilk@physik.uni-muenchen.de,
[9]jochen.kuhn@lmu.de, [10]nikola@merkas.tech, [11]s.kuechemann@physik.uni-muenchen.de

***Abstract*—As quantum technologies gain relevance across scientific and industrial domains, accessible educational frameworks are critical to preparing the next generation of quantum-literate learners. Traditional instruction often fails to engage diverse audiences or convey abstract quantum concepts intuitively. This paper introduces a novel AI-assisted learning architecture centered on the Embodied Language Model (ELM) – a hybrid approach that integrates Reinforcement Learning (RL) and Large Language Models (LLMs) – to address key challenges in quantum education through adaptive, game-based learning.**

We evaluate this framework in *Qookies*, a story-driven point-and-click adventure featuring the AI-controlled non-player character (NPC) Yuki as co-learner. Yuki combines an RL-based action model with an LLM which uniquely, like the player, begins with limited domain knowledge, acquiring understanding incrementally through shared game-play. The interaction design emphasizes observational learning, instruction, and dialogue: the player prompts Yuki to act, requests assistance, or engages in open conversation, while Yuki defers to the player when uncertain, suggests interactions with objects outside her reach, or explains quantum concepts contextually.

Key contributions include the RL model's runtime learning of object concepts through game-play observation, the LLM's prompt accumulation by narrative progression and user dialogue, and the interface design enabling communication between RL and LLM models, collectively mirroring the player's learning progress. Our architecture integrates these complementary processes to enable adaptive, personalized learning through collaborative exploration.

Evaluation results from related studies suggest that the game enhances learners' conceptual understanding of quantum phenomena, whereas the grounding RL component effectively reduces intrinsic cognitive load, simplifying the acquisition of complex concepts and promoting lasting learning through contextualized game-play. Overall, our framework offers a scalable, adaptive model for AI-assisted personalized education, contributes to hybrid AI architectures in educational technology, and suggests potential for domain transfer across STEM learning environments.

***Index Terms*—Game-based Learning, AI-assisted Education, Reinforcement Learning, Large Language Models, Personalized Learning, Educational Technology, Embodied Language Model**

I. Introduction and Background

There is a rapidly growing need for broad-based quantum literacy across science and education, as quantum technologies move from laboratories into industry and engineering workflows. Quantum concepts – superposition, entanglement, measurement – are abstract and counterintuitive, which poses a substantial barrier for learners without a strong mathematical background. Traditional lecture-based approaches often fail to provide the embodied intuition or contextualized experience learners need to form robust and sustainable mental models of these ideas.

Game-based learning offers a promising path forward by embedding abstract concepts in interactive tasks and narratives that provide situated, motivational contexts for exploration. Games can encourage to engage with complex topics in the first place and scaffold complex ideas through progressive, interactive challenges while providing feedback continuously; they thereby support active learning and reflection [1]. At the same time, modern AI systems (particularly LLMs and RL agents) enable novel adaptive and social forms of support – e.g., on-demand explanations, model demonstrations, and co-play – that can personalize the learning trajectory in real time. Thus, integrating game mechanics with AI affords an opportunity to design educational experiences that are simultaneously motivating, scaffolded, and adaptive.

When designing complex learning content, considerations regarding cognitive load are of crucial importance. Cognitive Load Theory argues that instructional designs must manage

This research was funded by the German Ministry for Research, Technology and Space

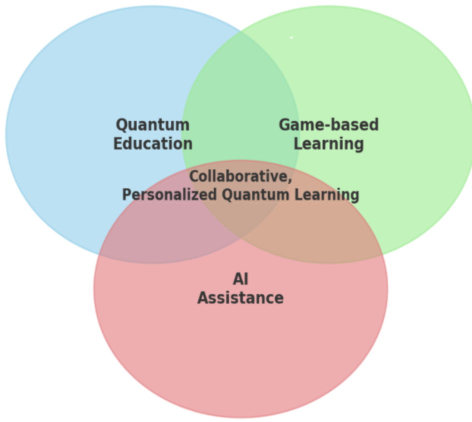

Fig. 1. Conceptual positioning of the GALaQSci project, illustrating the convergence of quantum education, game-based learning, and AI-assisted co-learning through the NPC Yuki.

intrinsic load (inherent difficulty of the subject matter), reduce extraneous load (unnecessary cognitive burden due to poor instructional design), and support germane load (effort required for schema construction and automation) [2]. Measuring and optimizing these loads is essential when introducing AI-mediated supports and new interaction patterns so that AI assistance reduces rather than adds to learners' processing burden. Empirical instruments for measuring intrinsic, extraneous and germane load are available and widely used in educational research.

A. The GALaQSci Project

The aim of the GALaQSci (Game-based and AI-assisted learning about Quantum Science) project is to develop a game that offers a low-threshold introduction to the subject of quantum physics and quantum technologies. Previously, Seskir et al.[3] highlighted the value of quantum technology games for outreach and educational purposes. One requirement for the level design in GALaQSci is that the puzzles should be solvable without any prior knowledge of physics or mathematics. At the same time, the game is intended to appeal to a broad target group, and should also be interesting for people who may already have some background knowledge and are interested in learning more. The addition of an AI-controlled non-player character (NPC) named 'Yuki' is motivated by the ambition to enable a collaborative and personalized learning experience in which the player receives appropriate support. Rather than an all-knowing AI, Yuki is designed as a co-player with whom a game world of quantum phenomena can be discovered in a playful manner. Taken together, GALaQSci is at the intersection of quantum education, game-based learning and AI assistance to promote collaborative, personalized learning (see Fig. 1).

B. Related Work

Yannakakis and Togelius [4] provide a comprehensive foundation on artificial intelligence methods in games, emphasizing how adaptive agents, procedural content generation, and player modeling can foster engagement and learning, principles that directly inform the design of our AI-assisted educational framework.

Schrodt outlines neurocomputational mechanisms of action understanding based on predictive coding and embodied simulation, offering a theoretical foundation for the co-learning and embodied cognition principles that underpin our approach [5]. The work also demonstrates how distinct (neural) modules with specialized functions can recursively communicate and synchronize to represent actions across multiple modalities, supporting coherent learning and embodied action simulation.

Building on the cognitive modeling paradigm, Schrodt et al. demonstrated that games (i.e., Super Mario) can serve as an effective testbed for studying adaptive cognitive architectures and social interaction in embodied agents. Their work introduced RL models capable of event-based reasoning, cooperative problem solving, and incremental learning through dialogue with the player as well as autonomous or instructed game-play [5], [6], [7], [8]. Moreover, the approach presented here also incorporates the relevant content- and function-based principles of collaborative knowledge construction pointed out by Fischer et al. (2002) [9]: (1) externalization and elicitation of task-relevant knowledge, as well as (2) conflict- and integration-oriented consensus building. To account for this, the AI model learns in parallel to the user by their actions and is able to communicate the task-relevant knowledge with the aim to find a joint solution.

Building on this line of research, our proposed Embodied Language Model (ELM) extends such cognitively inspired architectures by integrating Reinforcement Learning and Large Language Models into a unified, interactive framework.

II. AI-Assisted Learning Architecture

This section presents the conceptual architecture underlying the AI-assisted learning framework used in *Qookies*. It details how Reinforcement Learning and Large Language Models are integrated into a unified Embodied Language Model (ELM) that supports adaptive, grounded co-learning between the player and the NPC.

Note that while major components of the ELM have been implemented and evaluated within the *Qookies* environment, the architecture depicted here represents a full, theoretical design. It illustrates the complete set of components, interactions, and capabilities envisioned for the system, serving as a guiding blueprint for ongoing development and refinement in this and other applications.

A. Requirements for Enabling Co-learning Experience

To enable a co-learning experience, Yuki must start with limited domain knowledge so that its learning trajectory is parallel to the player's. Joint learning is particularly effective here, as it partially relieves the player of the burden of remembering what they have already learned, allowing them to concentrate on new content. At the same time, it does not raise the expectation that the AI knows everything, which may hinder independent and active thinking. The AI is designed to acquire knowledge incrementally via three channels

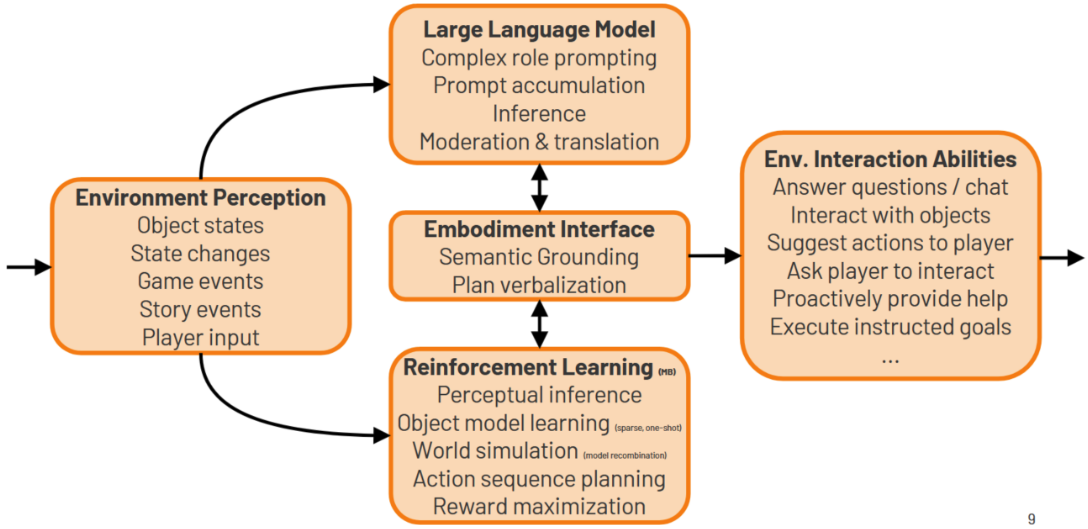

Fig. 2. ELM architecture flowchart. The Environment Perception module tracks object states, game events, and player inputs to maintain the AI's situational awareness. Perception is processed by both the LLM and RL modules by collecting prompts and learning object models. The embodiment interface mediates between RL and LLM representations to enable a range of environmental interaction capabilities for the AI.

– observation of player actions, active interaction with the environment (instructed, proactive or explorative), and natural-language dialogue. Crucially, the AI's actions and explanations must remain tied to the game's affordances to effectively avoid hallucinations and ensure consistent, contextually grounded behavior. These requirements – novice initial state, incremental learning, and tight perceptual-action grounding – motivate the modular AI architecture described next.

B. Concept and System Overview

We propose a hybrid architecture that couples a large language model (LLM) with a model-based Reinforcement Learning (MB-RL) controller through an embodiment interface to form an *Embodied Language Model* (ELM, see Figure 2). The LLM provides flexible natural language understanding and generation (explanations, dialog, instruction parsing, and prompt accumulation), while the RL component supplies situated, sensorimotor interaction and online concept acquisition through observation and action. The embodiment interface mediates between symbolic/linguistic representations (prompts, dialogue state) and grounded, game-state observations (object states, affordances, reachable actions). This bidirectional grounding is intended to reduce hallucination by constraining language outputs to the current embodied context, while preserving the LLM's strengths for explanation, narrative continuity, and creativity.

C. Architecture Modules, Interactions and Capabilities

The architecture is organized into the following modules:

Environment Perception: The Environment Perception module continuously monitors and encodes the current state of the game world, including object states, state changes, game and story events, and player inputs. It provides both AI modules – the LLM and RL components – with an up-to-date situational representation that grounds decision-making, dialogue, and action planning in the observable game world.

LLM Module: The LLM is responsible for turn-taking in dialogue, explanation generation, and high-level planning suggestions. In our current prototype of *Qookies* we use a state-of-the-art instruction-tuned model (Llama 3.1 70B [10]). The model is strongly prompted for its role in the narrative of the game – Yuki stays very much on the topic of Quantum Physics. Furthermore, in order to provide suitable help, the prompt contains a description about the current tasks. The LLM is also prompted about the currently present objects and changes in their state to increase the situational awareness. All story dialogues and conversation with the player, as well as game events, are accumulated in the prompt to preserve narrative continuity across levels, and also across sessions. Besides the strong prompt itself, a separate moderation system filters out harmful inputs. The model is also set up for supporting multiple languages.

MB-RL Module: The MB-RL module learns object-centered predictive models by decomposing the environment perception (perceptual inference). These decomposed models are learned and integrated via one-shot learning, enabling fast knowledge acquisition. Perceptual generalization can be achieved via predefined or neural mini-generalizers. Further, the models are sparsely connected and structured in a way that allows asymptotically efficient re-combination. This enables the simulation of and planning towards partial world states beyond the prior observations, enabling combinatory, one-shot generalization. A sequence of actions is then selected based on the maximum overall outcome in these partially simulated world states. However, real-time learning from scratch naturally means that the RL model may fail when creating an action plan. This is entirely intentional and results from the fact that both the AI and the player operate on an incomplete level of experience. However, classic RL theory and practice guide this component's design: MB-RL affords sample efficiency and the capacity to plan under uncertainty, which are important to enable a co-learning experience, where e.g. pre-training of a model-free approach would not be feasible.

Embodiment Interface: The bidirectional embodiment interface serves two complementary roles. First, it converts MB-RL action plans – sequences of actions executed on specific object identities – into an intermediate, structured representation that the LLM can interpret and render in natural language. Second, it translates language inputs from the user or other game characters into goals, state updates, or perceptual information that can be processed by the MB-RL module, effectively grounding high-level instructions in the game environment.

If MB RL is unable to generate a valid action plan, the LLM can produce purely prompt-driven replies, which may include more creative suggestions but are less tightly coupled to the game mechanics. In some cases, MB-RL failure is intentional and integrated into game-play as a design feature.

Environmental Interaction Capabilities: This module enables the AI to actively engage with both the player and the game world by translating internal representations into concrete behaviors. It allows the NPC to answer questions and chat, manipulate objects directly, suggest actions to the player, prompt the player to interact with otherwise inaccessible elements, proactively provide guidance, and execute instructed goals. By integrating outputs from both the LLM and MB-RL modules through the embodiment interface, this component operationalizes the co-learning experience, ensuring that the AI's actions and communications remain contextually grounded, responsive, and pedagogically effective.

The game- and learning environment *Qookies* and the interaction patterns resulting from this AI design are described in the following Sections.

III. Game-Based Learning Environment: Qookies – A Quantum Quest

This section introduces the *Qookies* game environment, where the ELM architecture is embedded and evaluated. It explains how the game's narrative, puzzle mechanics, and AI-driven interactions jointly operationalize quantum concepts through hands-on, story-based learning.

A. Educational and narrative design

Qookies – A Quantum Quest is a story-driven point-and-click adventure that embeds quantum learning objectives in puzzle mechanics and narrative progression. Each level combines escape-room style challenges with small, scaffolded experiments that operationalize key quantum concepts (e.g., qubits, superposition, entanglement) as well as illustrative phenomena such as polarization and fluorescence via manipulable game elements. The design targets secondary-school and early-university learners with levels that progress in conceptual complexity while supporting diverse entry points through optional hints and demonstrable NPC actions.

Qookies is being developed for mobile platforms (Android and iOS) with an intended release window in 2026, and is part of the GALaQSci initiative to broaden access to quantum topics through playful, AI-assisted experiences.

In *Qookies*, the AI-co-player Yuki is controlled by our ELM model, and is designed to inhabit an epistemic state that parallels a novice learner: limited initial knowledge, explicit uncertainty markers, and capacity for both action and language. This deliberate symmetry, with the AI-NPC acting as a co-learner rather than a teacher, is intended to create pedagogical situations in which players engage in joint problem-solving and reflection. The NPC's behavior provides a visible model of learning dynamics, allowing players to observe, compare, and adapt strategies collaboratively within the game environment.

Note that the ELM architecture just as the *Qookies* game itself are currently not fully implemented. Thus, not all of the capabilities of the theoretical design shown in Section II are evaluated here.

B. Player–NPC Interaction Patterns and Learning Dynamics

The interactions between the player and Yuki illustrate the collaborative learning dynamics supported by the AI-NPC. In Figure 3, the player directs Yuki to perform the next steps in the level, triggering RL-based inference that results in concrete actions on objects, such as placing levers and setting their states, which progress the game and scaffold the introduction of quantum concepts. Figure 4 shows a free-form chat interaction, where Yuki responds to player questions while staying focused on quantum science, demonstrating the LLM's role in contextual dialogue. Figure 5 highlights the interplay between RL and LLM: when the player asks for a hint, RL proposes an action plan if possible, which is verbalized by the LLM; if RL cannot generate a plan, the LLM provides a creative, prompt-driven response, maintaining engagement. Finally, Figure 6 depicts a situation where the AI recognizes that the player must act to achieve a goal, verbalizing the plan via the LLM to coordinate joint problem-solving. Together, these examples illustrate how observation, dialogue, and guided action combine to support a co-learning experience.

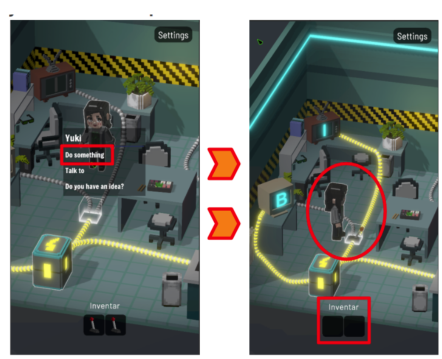

Fig. 3. Interaction example 1: The player asks Yuki to perform the next steps in the level. This triggers RL model inference, which in this case results in levers being taken from the inventory, placed in item slots, and their states (on/off) being set correctly. This causes the letters 'B', 'I' and 'T' to light up on monitors, completing the level and preparing the player for the next level, which is about Qubits and allegorically introduces levers that are on, off or both at the same time.

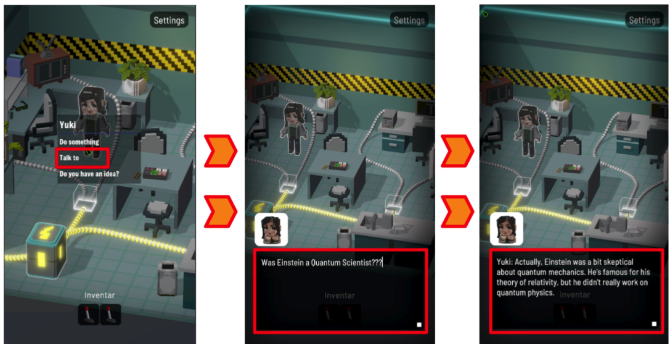

Fig. 4. Interaction example 2: The player is able to freely chat with Yuki about anything. Yuki is strongly prompted to discuss and stay on the topic of quantum science. When asked a general question about Einstein's relation to quantum mechanics, Yuki responds accordingly.

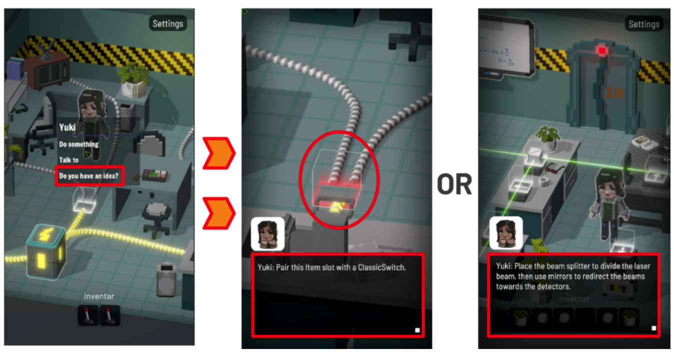

Fig. 5. Interaction example 3: The player asks for a hint (left), which triggers RL inference. If the RL model is able to find a rewarding sequence of actions based on its incomplete knowledge, the resulting action plan is converted by the Embodiment Interface into an intermediate language, which is then interpreted by the language model and converted into natural language (middle). If RL inference is unable to simulate rewards, the LLM inference is triggered directly and creates a more creative, but less grounded response based on the prompt history and level-dependent task description.

Fig. 6. Interaction example 4: In this situation, Yuki has to turn on a laser to solve the level. However, the laser is behind a locked door in another room, where the player is. Being asked to act in the current situation, RL infers a plan that involves actions of the player. Plan verbalization via the LLM results in Yuki asking the player to turn on the laser.

IV. Evaluations

To assess the educational impact of the proposed framework, an empirical study was conducted to examine how different AI configurations affect learning outcomes, cognitive load, and engagement in *Qookies*. This section summarizes the study design, expected learning mechanisms, and key results.

A. Expected Outcomes

Based on the game design and the ELM architecture, we anticipate a set of complementary educational benefits. Game-based learning, through narrative, puzzles, and progressive scaffolding, is expected to increase engagement and provide stronger contextual anchors for abstract quantum concepts. Shared and cooperative gameplay with an AI co-learner should foster social presence, attachment, and motivation, supporting reflection and collaborative strategy development. The co-learning design, in which both player and NPC acquire knowledge incrementally, aims to encourage active participation and deeper processing. Context-aware dialogue and prompt accumulation are intended to produce personalized, narrative-coherent support, while the tight integration of AI and game mechanics ensures that learning is closely coupled with play, minimizing risks of out-of-context information (hallucinations).

It is important to note that these anticipated effects span multiple cognitive, motivational, and social dimensions and cannot be fully investigated simultaneously. For the present study, we focus primarily on measures of conceptual understanding and cognitive load as broad, general indicators of the game's educational impact. These markers provide an initial, tractable means to assess whether the design supports meaningful learning outcomes, while leaving more detailed investigations of motivation, collaboration, and personalization for future work.

B. Results

To evaluate the educational impact of the AI design, Wermann et al. [11] (manuscript in preparation) conducted an extensive user study with over 150 participants, divided into three experimental groups. In the control group, participants interacted with the NPC Yuki only in a passive manner: Yuki provided background story and level objectives but offered no additional support. In the second group, participants could communicate with Yuki via text-based dialogue powered by a Large Language Model (LLM), allowing them to ask questions about puzzles, quantum concepts, or the NPC itself. Finally, in the third group, participants engaged with the Embodied Language Model (ELM), which combined the LLM dialogue capability with reinforcement-learning-driven demonstrations and assistance: Yuki could both chat and actively manipulate game objects to illustrate solutions, providing multimodal, co-learning support that mirrored the player's problem-solving and learning process.

The analyses indicate a robust game effect: There was a significant gain in conceptual understanding scores in a questionnaire that was completed once before and once after playing *Qookies* across all groups, which is consistent with learning through game-play.

Intrinsic Cognitive Load (ICL) reflects the mental effort required to process complex information, particularly when multiple interdependent elements must be integrated. Wermann et al.'s results demonstrate that learning with the Embodied Language Model (ELM) is measurably easier than with a text-only LLM: participants in the ELM condition experienced a significant reduction in ICL compared to those interacting with the LLM alone. This finding highlights the educational advantage of grounding language models in embodied, actionable experiences, showing that coupling dialogue with demonstrable interaction can make complex concepts more cognitively accessible and reduce barriers to effective learning.

Further methodological details, the complete analysis, and full results will be provided in the forthcoming article [11].

V. Discussion and Future Work

Quantum literacy is an urgent educational priority as quantum technologies move from research labs into classrooms and industry. Because of the abstract and counterintuitive nature of core quantum concepts, effective teaching must ground ideas in manipulable, perceptually rich experience. Combining game-based learning with adaptive AI support offers a promising route to make these ideas accessible to younger learners and to spark sustained interest in STEM. Our hybrid Embodied Language Model (ELM) leverages the complementary strengths of LLMs (flexible, contextualized dialogue; narrative continuity) and MB-RL (situated planning; demonstrable action) to create a co-learning partner that models and supports novice reasoning.

Early empirical signals indicate that the ELM design produces measurable pedagogical benefits compared with text-only LLM support [11]. In particular, the ELM condition showed a statistically significant reduction in intrinsic cognitive load relative to an LLM-only condition, while conceptual-understanding gains were observed across groups. These outcomes suggest that coupling dialogue with embodied, demonstrable interaction makes complex material more cognitively accessible. Mechanistically, we argue this effect arises from two interacting properties of the architecture: embodied demonstrations and tightly grounded language. Demonstrations externalize procedural steps and perceptual constraints, thereby reducing working-memory demands during problem solving; concurrently, the embodiment interface constrains LLM outputs to current affordances and structured MB-RL observations, which reduces out-of-context or speculative language. Together these mechanisms both ease moment-to-moment processing and improve the pedagogical relevance of AI-generated explanations.

Beyond these immediate learning gains, the architecture addresses a major challenge in educational AI – reliability. By translating MB-RL action plans into structured intermediate representations and returning grounded perceptual summaries to the LLM, the embodiment interface effectively anchors generative output in the game's current affordances and state, ensuring that explanations remain accurate, relevant, and verifiable.

While these findings are encouraging, they represent early signals. The current prototype of *Qookies* focuses on broad outcome markers such as conceptual understanding and cognitive load; disentangling effects on motivation, engagement, and reflection will require targeted follow-up studies.

Future work will extend this foundation in several directions: (1) studying how narrative framing and storytelling influence learning and engagement; (2) developing an AR mode to enhance embodied interaction; (3) integrating more proactive and context-sensitive AI behaviors for smoother gameplay support; and (4) adapting the ELM design pattern to other STEM domains such as chemistry labs, circuits, and systems biology.

References

[1] D. G. Oblinger, "The next generation of educational engagement," *Journal of interactive media in education*, vol. 2004, no. 1, pp. 10–10, 2004.

[2] J. Sweller, "Cognitive load during problem solving: Effects on learning," *Cognitive Science*, vol. 12, no. 2, pp. 257–285, 1988, Classic Cognitive Load Theory paper. DOI: 10.1207/s15516709cog1202_4.

[3] Z. C. Seskir et al., "Quantum games and interactive tools for quantum technologies outreach and education," *Optical Engineering*, vol. 61, no. 8, pp. 081 809–081 809, 2022.

[4] G. N. Yannakakis and J. Togelius, *Artificial Intelligence and Games*. Cham: Springer, 2018, Textbook on AI techniques for games.

[5] F. Schrodt, "Neurocomputational principles of action understanding: Perceptual inference, predictive coding, and embodied simulation," Ph.D. dissertation, Dissertation, Tübingen, Universität Tübingen, 2018.

[6] F. Schrodt, J. Lohmann, and M. V. Butz, "Mario becomes social," in *Video Proceedings of the 30th Conference of the Association for the Advancement of Artificial Intelligence*, 2016.

[7] F. Schrodt, Y. Röhm, and M. V. Butz, "An event-schematic, cooperative, cognitive architecture plays super mario," *Cognitive Robot Architectures*, vol. 10, pp. 10–15, 2017.

[8] F. Schrodt, J. Kneissler, S. Ehrenfeld, and M. V. Butz, "Mario becomes cognitive," *Topics in cognitive science*, vol. 9, no. 2, pp. 343–373, 2017.

[9] F. Fischer, J. Bruhn, C. Gräsel, and H. Mandl, "Fostering collaborative knowledge construction with visualization tools," *Learning and instruction*, vol. 12, no. 2, pp. 213–232, 2002.

[10] Meta AI, *Introducing llama 3.1: Our most capable models to date*, https://ai.meta.com/blog/meta-llama-3-1/, Blog / model announcement. Accessed: 2025-10-14, Jul. 2024.

[11] C. Wermann et al., "Interaction with ai-controlled non-player-characters in serious games," unpublished, 2025.