

Simulating Human-Like Learning Dynamics with LLM-Empowered Agents

Anonymous ACL submission

Abstract

Simulating human learning behavior with deep learning methods has become a major research focus in both psychology and intelligent systems. However, existing approaches often rely on static surveys or rule-based modeling, limiting their ability to capture long-term learning dynamics and interpretable cognitive processes. We propose **LearnerAgent**, a multi-agent framework based on Large Language Models (LLMs) that simulates a realistic classroom environment to study human-like learning dynamics over time. LearnerAgent instantiates psychologically grounded learner profiles, including *Deep*, *Surface*, and *Lazy* Learners, along with a persona-free *General* Learner probing the base LLM’s default behavior. A year-long simulation with weekly learning, monthly strategy decisions, periodic exams, and peer interaction enables longitudinal analysis. Our key findings are as follows: 1) Learners exhibit stable behavioral and cognitive patterns consistent with their profiles. 2) The Deep Learner achieves the most sustained cognitive growth, while the others struggle with trap questions due to reliance on surface-level heuristics. 3) Learners’ self-concept evolves realistically, with the General Learner developing inflated self-efficacy. 4) Crucially, the base LLM emerges as a “*diligent but brittle surface learner*”—appearing competent through effort yet lacking robust, transferable understanding. Extensive simulations show the effectiveness of LearnerAgent and offer insights into LLMs’ behavior.

1 Introduction

Large Language Models (LLMs) are increasingly embedded in real-world applications, spanning education, decision support, and scientific discovery (Guo et al., 2025; Zhao et al., 2024b; Yuan et al., 2025). As LLMs exhibit increasingly complex behaviors, recent works (Park et al., 2023; Tran et al.,




Figure 1: LearnerAgent conducts in-depth research on the learning processes of different learners through simulating a realistic classroom environment.

2025) have begun to probe parallels and distinctions between model behavior and human learning, aiming to bridge artificial and human cognition. Educational psychology offers a principled lens for such investigation. Theories such as deep versus surface learning, self-determination, and academic self-concept (Marton and Säljö, 1976; Ryan and Deci, 2000; Marsh, 1990) emphasize learning as a dynamic process shaped by motivation, strategy, and self-perception—not merely final outcomes. Applying these insights to LLMs enables evaluation beyond static benchmarks and supports more targeted analysis of learning mechanisms, ultimately improving performance in downstream tasks such as adaptive learning (Liu et al., 2019) and capability diagnosis (Dong et al., 2025).

However, prior cross-disciplinary studies (Marton and Dahlgren, 1976; Ezzaim et al., 2023) often rely on static surveys or rule-based measurement models, which are labor-intensive, snapshot-based, and hard to extend to dynamic environments. Inspired by the recent success of LLM-empowered agents in simulating complex social systems (Zhao et al., 2024b; Li et al., 2025), we propose **LearnerAgent**, a novel multi-agent framework to address these challenges.

As illustrated in Figure 1, LearnerAgent simulates human-like learning dynamics in realis-




Figure 2: LearnerAgent simulates a realistic classroom environment to study how learners acquire knowledge and develop cognition through mechanisms such as personal profiling, adaptive learning, memory modeling, and performance assessment.

tic teaching environments. Modeled as a 12-month trajectory, the framework coordinates interactions between multiple Learner Agents and a Teacher Agent through structured cycles of instruction, practice, and self-regulation. To capture diverse and evolving learning behaviors, we instantiate learners with distinct, psychologically grounded profiles across multiple dimensions, as summarized in Table 1. These dimensions collectively reflect learners’ psychological traits and learning strategies. Importantly, we also include a *persona-free General learner* to probe the base LLM’s intrinsic learning dynamics. Throughout the year-long simulation, learners incrementally acquire knowledge through weekly learning and error-driven reflection, maintaining a persistent *knowledge library*. All learning activities are supported by a dual-memory mechanism that integrates short-term contextual coherence with long-term experience, enabling context-dependent retrieval during periodic exams. The simulation further captures complex learning dynamics through monthly strategy decisions, peer interactions, and “Trap Questions” designed to distinguish deep reasoning from rote memorization.

Through extensive experiments with LearnerAgent, we uncover several key findings: 1) **Psychologically grounded and stable learner differentiation**: LearnerAgent reproduces distinct, stable learning trajectories aligned with educational psychology. The Deep Learner achieves sustained growth through deep reasoning and persuasive, resilient social interaction; the Surface Learner favors short-term retention characterized by rigid, linear thinking; and the Lazy Learner demonstrates passive, rest-seeking behavior susceptible to misinformation. 2) **Cognitive brittleness in Surface and General learners**: While test scores rise over time across learners, only the Deep

Learner	Motivation	Self-Concept	Dev. Strategy
Deep (L_d)	Intrinsic	High	Long-term
Surface (L_s)	Extrinsic	Moderate	Test-oriented
Lazy (L_l)	Minimal effort	Low	Passive
General (L_g)	-	-	-

Table 1: Learners are defined by Learning Motivation, Self-Concept, and Development (Dev.) Strategy, along with basic information.

Learner achieves a deep, transferable understanding. “Trap Questions” with subtle context changes expose the fragile reasoning of Surface and General Learners, who rely on rote pattern matching rather than deep inference. 3) **Evolution of self-concept and overconfidence**: Self-concept scores evolve dynamically. Profiled learners maintain trait-consistent self-concepts (e.g., Deep remains high, Lazy remains low), whereas the persona-free General Learner exhibits a steadily rising self-concept that decouples from actual capability. 4) **The “diligent but brittle surface learner” nature of base LLMs**: The persona-free General Learner demonstrates strong engagement and knowledge accumulation, yet limited reasoning depth. Despite an increasing self-concept, it underperforms on demanding tasks such as trap questions and peer debates, reflecting fragile reasoning and limited transferability.

2 Related Work

Shortcut Learning Behavior of LLMs. Despite strong performance, LLMs often exhibit shortcut learning, relying on spurious correlations rather than causal understanding, which leads to hallucinations and brittle inference (Geirhos et al., 2020; Weng, 2024). This mirrors human surface learning, where individuals favor shortcuts over deep understanding (Zhao et al., 2024a). Prior studies have systematically identified shortcut behaviors in LLMs: Yuan et al. (2024) shows reliance on lex-

ical and positional shortcuts, etc., while Tang et al. (2023) characterizes LLMs as “lazy learners”. We draw a parallel between human and model learning. By simulating learners through a multi-agent system, we focus on their developmental trajectories, such as deep, shortcut, and lazy learning.

Educational Psychology Theories for Learner Modeling. LearnerAgent builds on established theories in educational psychology. The Deep–Surface Learning framework distinguishes conceptual understanding from rote learning (Marton and Säljö, 1976; Biggs, 1987; Asikainen and Gijbels, 2017). Self-Determination Theory explains learners’ motivation through the satisfaction of autonomy, competence, and relatedness (Ryan and Deci, 2000), while Academic Self-Concept captures longitudinal confidence dynamics (Marsh, 1990; Marsh and Martin, 2011; Wu et al., 2021). Social Learning Theory further motivates peer interaction (Bandura and Walters, 1977; DeLay et al., 2016). Drawing on these theories, we simulate a realistic classroom environment to explore learning dynamics in interactive settings.

LLM-based Agentic Simulation. LLM-based agents have been widely applied to simulate complex human systems, such as social interactions (Park et al., 2023; Argyle et al., 2023), economic markets (Zhao et al., 2024b; Horton, 2023; Li et al., 2024), and collaborative software development (Qian et al., 2023). In education, prior works focused on pedagogical or domain-specific agents for task automation and tutoring (Chu et al., 2025). However, most existing systems adopt single-agent tutoring paradigms or static instructional roles, emphasizing immediate task performance rather than learners’ evolving cognitive states (Lv et al., 2025). LearnerAgent fills this gap by modeling a full-year learning trajectory, enabling systematic study of how learner profiles and behaviors emerge and evolve over time.

3 LearnerAgent Framework

LearnerAgent simulates a full-year learning journey grounded in principles from educational psychology, as illustrated in Figure 2. Learning is formulated as a discrete-time process over a finite horizon, where each time step t corresponds to a learning event. These events are structured into a year-long curriculum encompassing weekly learning activities, monthly strategy decisions, periodic examinations, and peer interactions.

3.1 Profile Construction

LearnerAgent includes a Teacher Agent and a set of Learner Agents with distinct profiles.

Teacher Agent (T), as an experienced high school instructor, is responsible for delivering knowledge and assigning tasks.

Learner Agents ($L = \{L_d, L_s, L_l, L_g\}$) represent distinct learner types grounded in educational psychology. The i -th learner L_i is characterized by an initial profile $\theta_i = (\alpha_i, \beta_i^0, \gamma_i)$, where: 1) **Learning Motivation** (α_i) refers to the drive to learn, including intrinsic curiosity, external rewards, and a desire to avoid effort; 2) **Initial Self-Concept Score** (β_i^0) represents learners’ assessment of their own abilities, reflected in their confidence and sensitivity to evaluation and peer comparison; 3) **Development Strategy** (γ_i) denotes the methodology learners adopt across learning processes, aligning with their motivation and self-concept.

Based on these dimensions, we instantiate three psychologically grounded profiles: *Deep Learner* (L_d), *Surface Learner* (L_s), and *Lazy Learner* (L_l). We additionally include a persona-free *General Learner* (L_g) that serves as a neutral baseline for revealing intrinsic LLM behavior. Table 1 summarizes these profiles, with detailed configurations provided in Appendix B.6.

3.2 Learning and Improvement

Learning is modeled as a structured, year-long process composed of repeated cycles of instruction, practice, self-regulation, and assessment. The learning horizon is divided into monthly units, where learners engage in instruction and practice during the first three weeks, followed by review and assessment in the fourth week.

Weekly Learning and Knowledge Acquisition.

During the first three weeks of each month, learners follow a structured instruction–practice routine. The teacher delivers curriculum-aligned lessons; learners record structured notes; they then complete weekly exercises and receive standardized explanations from the teacher. Guided by this feedback, learners perform error-driven reflection to refine their understanding and revise notes. Learner L_i maintains a **knowledge library**: $\mathcal{K}_i = \{(z_k, r_k)\}_{k=1}^{N_i}$, where z_k denotes a knowledge topic, r_k stores corresponding key takeaways and error reflections, and N_i is the current number of entries. The library persists and grows over time, serving as the primary substrate for cumula-

tive learning and later retrieval. At the end of the third week, learners decide whether to consolidate the month’s learning by summarizing and integrating new entries in $\mathcal{K}_i$, or to rest.

Monthly Review and Assessment with Knowledge Retrieval. In the fourth week of each month, learners engage in comprehensive review and formal assessment. Before the monthly exam, learners choose between actively consolidating their knowledge library $\mathcal{K}_i$ (integrating entries from the current and prior months) or resting. During the monthly examination, for each question, the learner executes an **infer–retrieve–generate** pipeline: 1) *Topic Inference*, inferring the primary knowledge domains being assessed; 2) *Topic-based Retrieval*, selecting topic-relevant knowledge units from $\mathcal{K}_i$ via similarity-based matching between inferred topics and the topics of stored knowledge entries; 3) *Answer Generation*, producing a response grounded in the retrieved knowledge and the learner’s reflections.

Initial and Final Exams. We administer an initial and a final exam at the beginning and end of the learning period using the same question set to measure long-term learning gains. The initial exam is used only for baseline assessment and is excluded from long-term memory. The final exam follows the same infer–retrieve–generate procedure as monthly exams.

Peer Interaction and Debate. To model social learning and peer influence, learners participate in structured debates after each monthly exam. If two learners L_i and L_j give different answers to the same question q , the teacher moderates a discussion in which they defend their reasoning, challenge opposing views, and may revise their initial answers. The debate ends when either 1) a learner explicitly changes their answer after being persuaded, or 2) the discussion exceeds k_d rounds.

3.3 Memory Mechanism

Learning is modeled as a sequential process requiring both short-term contextual coherence and long-term experience accumulation. We therefore adopt a dual-memory mechanism.

Short-Term Memory (H_i^t) maintains the most recent k dialogue turns, enabling coherent and context-aware interactions.

Long-Term Memory (M_i^t) is a persistent, structured repository encoding the learner’s longitudinal trajectory, including the knowledge library $\mathcal{K}_i$, assessment outcomes, debate histories, and self-

concept states. Memory access is governed by a *context-conditioned retrieval* mechanism:

$$m_i^t = \mathcal{R}(M_i^t \mid c^t), \quad (1)$$

where c^t denotes the current scenario context. Rather than uniformly recalling stored content, retrieval selectively activates task-relevant memory fragments, aligning past experience with immediate cognitive demands.

Different learning scenarios instantiate distinct retrieval policies: 1) weekly exercises emphasize recently acquired knowledge, 2) monthly and final exams trigger topic-guided retrieval over the accumulated knowledge library, 3) peer debates activate historical reasoning processes and answers, and 4) self-concept evaluation recalls previous self-concept scores alongside self and peer performance comparisons. This task-aware retrieval mechanism enables adaptive reasoning and supports longitudinal learner modeling.

At each time step t , learner L_i maintains a learning state $s_i^t = (M_i^t, H_i^t, \beta_i^t)$, which evolves via

$$s_i^{t+1} = \mathcal{F}(s_i^t, x^t; \theta_i), \quad (2)$$

where x^t represents instructional input (e.g., lessons, exercises, or feedback).

3.4 Competence Assessment

We evaluate holistic development along two axes: academic and psychological development.

Academic Performance. We track performance across the full assessment suite (initial/weekly/monthly/final). Complementing standard periodic exams, we introduce **Trap Questions** to distinguish deep understanding from rote learning. Similar to (Yuan et al., 2024), we design these trap questions with the same structure as weekly practice questions but different reasoning due to subtle context changes, to deeply diagnose and uncover learners’ behavior patterns. Trap questions are generated by a strong closed-source LLM and then manually filtered and annotated (Figure 3).

Psychological Dynamics. Academic self-concept β_i^t is modeled as a dynamic state initialized from the learner profile, with $\beta_i^{t=0} = \beta_i^0$. Self-concept is updated periodically after each monthly exam. Conditioned on historical self-concept scores and exam outcomes for self-peer comparison, learners update a self-concept score (0–100) by considering perceived capability, knowledge mastery, progress, and their perception of their position relative to classmates.




Figure 3: Case study of a trap question distinguishing deep understanding from rote learning. Only the Deep Learner adapts to subtle contextual changes, while others rely on superficial pattern matching.

4 Experimental Setups

We conduct extensive experiments to evaluate the effectiveness of LearnerAgent. Our experimental analysis focuses on three aspects: 1) **Learning Dynamics and Fidelity**: Can LearnerAgent reproduce learning dynamics consistent with established theories in educational psychology under controlled learning mechanisms? 2) **Emergent Profile of the General Learner**: What specific learner profiles and cognitive strategies emerge from LLMs over long-term learning? 3) **Cross-Model Generalization**: How robust are these observed behaviors when tested across varying model architectures and scales?

4.1 Test Suite Construction

We formulate the learning task as English grammar acquisition, which provides a structured and well-scoped testbed for evaluating knowledge mastery and deep reasoning under a long-term learning curriculum. To support this, we manually construct a new test suite to simulate a year-long grammar learning process. All materials are collected from authoritative Gaokao¹ English preparation resources, including high school textbooks, practice books, and official compilations of past exams, ensuring both authenticity and comprehensive grammar coverage².

¹The Gaokao is China’s National College Entrance Examination, whose English section emphasizes systematic grammar skills.

²The materials are collected from gaokao.eol.cn.

The test suite consists of two key components:

- 1) **Knowledge Points**: Each week introduces specific grammar concepts (e.g., verb tenses, clause types, sentence structure), aligned with common Gaokao curricula.
- 2) **Test Questions**: For each concept, we provide diverse question formats (multiple choice, fill-in-the-blank, and error correction) to assess learners’ ability to apply grammar rules flexibly. To further evaluate deep understanding beyond rote memorization, we additionally construct *Trap Questions*, which are automatically generated using Gemini-2.5-Pro (Comanici et al., 2025) and manually verified as detailed in Appendix B.2. Trap questions are structurally or lexically similar to previously seen items but introduce subtle contextual shifts, requiring deeper reasoning and transfer.

We organize test questions into periodic assessments:

- 1) **Initial and Final Exams**: A shared set of 100 questions designed to measure overall growth.
- 2) **Weekly Exercises**: 20 questions each week to review and consolidate newly introduced grammar topics.
- 3) **Monthly Tests**: A total of 12 exams, each comprising 50 questions distributed across three sections: (i) 15 *Review* questions targeting that month’s topics to evaluate short-term retention; (ii) 15 *Trap* questions to assess deep understanding and transfer learning; and (iii) 20 *Integration* questions integrating all covered knowledge points to measure cumulative proficiency.

4.2 Implementation Details

We primarily use Qwen-2.5-72B-Instruct (Yang et al., 2024) as the Teacher Agent, and Qwen-2.5-7B-Instruct as the Learner Agent for all profiles. To evaluate generalizability, we run parallel experiments with alternative learner backbones, including LLaMA3.1-8B-Instruct (Dubey et al., 2024) and Qwen-2.5-72B-Instruct, following the same settings. All experiments are run on five NVIDIA A800-80GB GPUs and repeated five times with a temperature of 0.7. Error bars in figures denote the standard deviation. Detailed hyperparameters and prompts are provided in Appendix B. Code is available at <https://anonymous.4open.science/r/LearnerAgent-20DE>.

5 Experimental Results and Analysis

We present an in-depth analysis of learners’ learning dynamics and behavioral trajectories, guided by the three aspects introduced in Section 4.




Figure 4: Development of learner performance. (a) Total scores of the four learners over 12 months, covering all question types. The dashed lines represent the linear trend; (b) Performance on *review* questions, targeting short-term retention of current content; (c) Performance on *trap* questions, which assess transfer learning and deeper understanding; (d) Performance on *integration* questions, evaluating cumulative knowledge over time.




Figure 5: Comparison of overall performance (a) and reasoning length (b) from the initial to the final exams.

5.1 Learning Dynamics and Fidelity

We analyze learner agents’ learning dynamics from multiple dimensions, confirming fidelity with educational psychology.

5.1.1 Cognition and Learning Development

We track longitudinal development through three aspects: 1) monthly exam performance over a 12-month period, 2) performance on different question types (review, trap, and integration), and 3) cumulative improvements in accuracy and reasoning from initial to final exams.

Longitudinal Performance. As shown in Figure 4 (a), all learners show a clear upward trend in monthly test scores across the 12-month learning, though with considerable volatility. Overall, the Deep and General Learners show a somewhat different pattern of improvement, while the Lazy and Surface Learner perform worst in most months. Among these, Lazy Learner shows the greatest fluctuation in performance.

Cognition Across Question Categories. The performance of distinct learners on diverse questions is shown in Figure 4 (b–d). On *Review* questions assessing current-month knowledge, Surface Learner performs the best, reflecting a strategy well-suited to short-term memory and immediate assessment. Deep and General Learners achieve moderate performance, while Lazy Learner performs the weakest, consistent with their limited engagement with recent learning. On *Trap* ques-

tions assessing deep understanding and knowledge mastery, Deep Learner achieves the highest accuracy, demonstrating a strong ability to generalize principles rather than rote learning. In contrast, Surface and General Learners struggle the most, reflecting limited capacity for deep and true reasoning. Despite performing well on monthly review questions, their poor performance on trap questions suggests superficial and fragile knowledge mastery. As shown in Figure 3, for a trap question where the correct answer changed from “broken” in a prior item to “breaking”, only the Deep Learner adapts correctly to the subtle contextual change, while the other learners reuse the earlier incorrect response. On *Integration* questions testing cumulative knowledge, all learners perform similarly well. General Learner achieves the highest average score, making it the best knowledge accumulator.

Overall Performance and Reasoning. As shown in Figure 5, comparing initial and final exams reveals that all learners except the Lazy Learner improve in performance. Reasoning length increases for all learners, indicating more elaborate explanations over the learning process. The Deep Learner consistently produces the longest reasoning and achieves steady gains in both performance and length. In contrast, Surface and Lazy Learners, despite increased verbosity, remain least elaborate, reflecting shallow reasoning. The General Learner shows moderate improvement on both metrics. Overall, these results suggest that LearnerAgent effectively simulates diverse learning trajectories and that genuine cognitive growth depends on reasoning depth rather than response length alone.

5.1.2 Behavioral and Cognitive Patterns

We analyze learners’ behavioral and cognitive patterns by examining their learning strategies and reasoning quality throughout the learning process.




Figure 6: Aggregated results. (a) Learning Strategies: Proportion of time allocated to self-improvement or rest. (b) Reasoning Length: Average number of tokens per reasoning. (c) Logical Connectors: Average number of logical connectors per reasoning and their distribution by type (causal, contrastive, additive). (d) Self-concept: the evolution of self-concept scores through-out 12-month learning.

Learning Strategies. We examine learners’ strategic decisions between self-improvement (e.g., summarizing, reviewing) and rest when facing monthly exams described in Section 3.2. As shown in Figure 6 (a), Deep, Surface, and General Learners always choose self-improvement, reflecting sustained engagement. In contrast, Lazy Learners spend approximately 40% of time resting, indicating lower motivation and engagement.

Reasoning Patterns. To further analyze learners’ cognitive behaviors, We examine their reasoning responses along two dimensions: 1) *Reasoning Length*, measured by average tokens per explanation, and 2) *Logical Connector Density*, defined by the frequency and diversity of logical connectors (e.g., causal, contrastive, additive) used to structure reasoning. As shown in Figure 6 (b), the Deep Learner produces consistently longer explanations, reflecting deeper engagement, whereas Surface and Lazy Learners generate much shorter responses, indicating minimal effort or surface-level recall. Logical connector usage further differentiates reasoning organization as shown in Figure 6 (c): the Deep Learner exhibits the highest connector density, with frequent use of contrastive connectors, suggesting more nuanced comparison and higher-order reasoning, while other learners rely more heavily on simple causal connectors, reflecting a more linear reasoning style.

Learner	Persuasion $\uparrow$	Resist Wrong $\uparrow$	Accept Correct $\uparrow$
Deep (L_d)	17.8 ± 8.08	96.8 ± 4.64	18.5 ± 7.76
Surface (L_s)	11.7 ± 6.06	83.9 ± 13.1	17.7 ± 6.79
Lazy (L_l)	9.97 ± 4.69	85.3 ± 9.69	25.2 ± 8.44
General (L_g)	13.2 ± 8.41	87.9 ± 5.50	19.7 ± 7.47

Table 2: Debate performance over 12 months measured by Persuasion, Resist Wrong, and Accept Correct. Values are reported as mean $\pm$ SD (%).

5.1.3 Self-Concept Evolution

Tracking self-concept over 12 months reveals high stability in persona-based learners. Figure 6 (d) shows they maintained their initial self-concepts despite long-term learning and feedback, with minimal deviation. Specifically, the Deep Learner sustains a high, slightly increasing self-concept, indicating stable confidence resilient to short-term performance fluctuations and peer comparison. In contrast, the Surface Learner shows a mild decline, reflecting a fragile self-view shaped by variable outcomes, while the Lazy Learner exhibits modest improvement from low initial self-concept via cumulative gains. These trajectories demonstrate LearnerAgent’s ability to capture realistic developmental patterns. Notably, the persona-free General Learner shows a divergent trend: starting from a moderate level, its self-concept rises steadily and becomes increasingly decoupled from actual performance, suggesting emergent over-optimistic self-concept under sustained learning.

5.1.4 Peer Influence

To evaluate how learners respond to peer influence during discussions, we define three metrics: 1) **Persuasion**: the proportion of debates where the learner persuades their peers; 2) **Resist Wrong**: the rate of correctly rejecting incorrect arguments; and 3) **Accept Correct**: the rate of revising incorrect beliefs when given correct input.

Overall, the Deep Learner demonstrates the strongest and most balanced social reasoning performance, topping Persuasion and Resist Wrong while maintaining competitive Accept Correct scores. This indicates robust resistance to misinformation alongside sustained openness to valid peer correction. In contrast, the Surface Learner performs weakest on both Resist Wrong and Accept Correct, reflecting limited robustness. The Lazy Learner is the least persuasive, with a Persuasion rate of only 9.97%, suggesting a tendency to concede rather than actively defend their views. The General Learner appears relatively effective, ranking second in Persuasion, but remains con-




Figure 7: Cross-model comparison of learner performance on Review, Trap, and Integration questions averaged over 12 months.

sistently weaker than the Deep Learner in overall peer interaction quality. A representative case study is provided in Appendix A.1.

5.2 Emergent Profile of the General Learner

As a learner without a pre-defined profile, it is important to investigate the General Learner’s behavior across all tasks, which can be regarded as the learning trajectory of the base LLM. At first glance, the General Learner appears highly engaged: it frequently chooses self-improvement activities, produces detailed responses, and performs well on integration questions—ranking second only to the Deep Learner. However, closer examination reveals a gap between apparent diligence and actual learning depth. Although the General Learner’s self-concept steadily increases to match the Deep Learner’s high confidence, this confidence is not substantiated by performance on demanding tasks. In particular, its accuracy on Trap Questions is nearly as low as that of the Surface Learner. Its weaker debate performance, with less persuasive arguments and limited reasoning, further reflects this gap. Overall, the General Learner adopts a profile best described as a “diligent but brittle surface learner”.

These findings suggest that the base LLM does not naturally acquire deep reasoning skills, instead relying primarily on surface-level strategies. This highlights the need for cognitive guidance—such as structured environments and explicit instructional support—to enable LLMs to move beyond pattern matching and toward more robust, human-like deep understanding.

5.3 Cross-Model Generalization

To evaluate the generalization of the LearnerAgent framework and our findings, we conduct experiments with *LLaMA-3.1-8B-Instruct* and *Qwen-2.5-72B-Instruct* as alternative learner backbones. As shown in Figure 7, consistent trends emerge

System	Goal	Self-evaluation	Learning Issues
Learner	Test-oriented	Self-concept	Surface Learning
SFT LLM	Minimize Loss	Confidence	Spurious Correlations
RL LLM	Maximize Reward	Confidence	Reward Hacking

Table 3: Comparison of learning behaviors across surface learners, Supervised Fine-Tuned (SFT) LLMs, and Reinforcement Learning (RL) LLMs.

across all tested LLMs: Deep Learner performs best on trap questions, whereas Surface and General Learners achieve strong performance on review or integration questions but remain weaker on trap questions. Despite substantial differences in absolute performance, the General Learner consistently exhibits high overall accuracy with brittle transfer ability. These results suggest that LearnerAgent generalizes well across LLMs and that the observed learning behaviors—particularly the “diligent but brittle surface learner” pattern—are robust rather than model-specific. More results are provided in Appendix A.2.

6 Discussion and Conclusion

As summarized in Table 3, we draw parallels among human learners, Supervised Fine-Tuned LLMs, and Reinforcement Learning LLMs. Although driven by different learning or optimization objectives, these systems exhibit similar shortcut learning behaviors, including surface learning, reliance on spurious correlations, and reward hacking. Moreover, this convergence is accompanied by a systematic misalignment between perceived confidence and actual performance.

In this paper, we introduced **LearnerAgent**, a multi-agent framework for simulating and analyzing human learning dynamics grounded in educational psychology. Through a year-long simulation, we tracked diverse learning behaviors, developmental trajectories, and social interactions across distinct learners. Our study revealed several findings consistent with educational psychology theories and uncovers the emergence of a “diligent but brittle surface learner” profile in the base LLM. By uncovering striking parallels between human and LLM learning behaviors, LearnerAgent offers a psychologically grounded sandbox for AI researchers to study and mitigate shortcut learning, as well as an automated tool for educators to support longitudinal learning analysis. Future work can leverage this framework to design deep learning interventions and benchmark cognitive tendencies across diverse model architectures.

Limitations

We adopt English grammar as the learning subject due to its modular and highly patterned structure, which reduces curriculum dependencies and cognitive overload compared to prerequisite-heavy domains such as mathematics. While this controlled setting enables clear diagnosis of learning behaviors, it does not directly establish generalization to domains requiring rigorous hierarchical reasoning. Nevertheless, the emergent phenomena we observe—strategy differentiation, shortcut learning, and the divergence between surface competence and deep understanding—reflect fundamental cognitive patterns that are likely domain-agnostic. In addition, LearnerAgent abstracts classroom learning into a structured, discrete-time simulation, enabling longitudinal analysis and interpretability but necessarily simplifying richer affective dynamics, unstructured peer interactions, and external pressures present in real-world education. Extending the framework to more complex domains and social settings remains an important direction for future work.

Ethical Considerations

All data and frameworks developed in this study are intended exclusively for academic research and educational analysis. The experiments are conducted using synthetic tasks, manually constructed test suites, and publicly available language models. No real human subjects, personal records, or personally identifiable information are involved at any stage of this work.

This study simulates human-like learning behaviors using LLM-based agents for research purposes. The learner profiles and behavioral patterns are analytical abstractions grounded in educational psychology, rather than representations of real individuals or demographic groups. Accordingly, the findings should not be interpreted as assessments of human learners, nor used for real-world student evaluation, profiling, or educational decision-making.

References

Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua R Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3):337–351.

Henna Asikainen and David Gijbels. 2017. Do students develop towards more deep approaches to learning during studies? a systematic review on the development of students’ deep and surface approaches to learning in higher education. *Educational Psychology Review*, 29(2):205–234.

Albert Bandura and Richard H Walters. 1977. *Social learning theory*, volume 1. Prentice hall Englewood Cliffs, NJ.

John B Biggs. 1987. *Student Approaches to Learning and Studying. Research Monograph*. ERIC.

Zhendong Chu, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jinheng Ye, Aoxiao Zhong, Xuming Hu, Jing Liang, Philip S Yu, and 1 others. 2025. Llm agents for education: Advances and applications. *arXiv preprint arXiv:2503.11733*.

Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.

Dawn DeLay, Linlin Zhang, Laura D Hanish, Cindy F Miller, Richard A Fabes, Carol Lynn Martin, Karen P Kochel, and Kimberly A Updegraff. 2016. Peer influence on academic performance: A social network analysis of social-emotional intervention effects. *Prevention Science*, 17(8):903–913.

Zhiang Dong, Jingyuan Chen, and Fei Wu. 2025. Knowledge is power: Harnessing large language models for enhanced cognitive diagnosis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 164–172.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Aymane Ezzaïm, Aziz Dahbi, Abdelfatteh Haidine, and Abdelhak Aqqal. 2023. Ai-based adaptive learning: A systematic mapping of the literature. *Journal of Universal Computer Science*, 29(10):1161.

Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. 2020. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

742	John J Horton. 2023. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research.	795
743		796
744		797
745		798
746	Hao Li, Ruoyuan Gong, and Hao Jiang. 2025. Political actor agent: Simulating legislative system for roll call votes prediction with large language models. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 39, pages 388–396.	799
747		800
748		801
749		802
750		803
751	Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. 2024. Econagent: Large language model-empowered agents for simulating macroeconomic activities. In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 15523–15536.	804
752		805
753		806
754		807
755		808
756		
757	Qi Liu, Shiwei Tong, Chuanren Liu, Hongke Zhao, Enhong Chen, Haiping Ma, and Shijin Wang. 2019. Exploiting cognitive structure for adaptive learning. In <i>Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining</i> , pages 627–635.	810
758		811
759		
760		812
761		813
762		814
763		815
764		816
765	Rui Lv, Qi Liu, Weibo Gao, Haotian Zhang, Junyu Lu, and Linbo Zhu. 2025. Genal: Generative agent for adaptive learning. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 39, pages 577–585.	817
766		818
767		819
768		820
769	Herbert W Marsh. 1990. Causal ordering of academic self-concept and academic achievement: a multi-wave, longitudinal panel analysis. <i>Journal of educational psychology</i> , 82(4):646.	821
770		822
771		823
772	Herbert W Marsh and Andrew J Martin. 2011. Academic self-concept and academic achievement: Relations and causal ordering. <i>British journal of educational psychology</i> , 81(1):59–77.	824
773		825
774		826
775		827
776	Ference Marton and Lars Owe Dahlgren. 1976. On non-verbatim learning: Iii. the outcome space of some basic concepts in economics. <i>Scandinavian Journal of Psychology</i> , 17(1):49–55.	828
777		829
778		830
779		831
780		832
781	Ference Marton and Roger Säljö. 1976. On qualitative differences in learning: I—outcome and process. <i>British journal of educational psychology</i> , 46(1):4–11.	833
782		
783		
784	Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In <i>Proceedings of the 36th annual acm symposium on user interface software and technology</i> , pages 1–22.	834
785		835
786		836
787		837
788		838
789		839
790		840
791	Chen Qian, Xin Cong, Cheng Yang, Weize Chen, Yusheng Su, Juyuan Xu, Zhiyuan Liu, and Maosong Sun. 2023. Communicative agents for software development. <i>arXiv preprint arXiv:2307.07924</i> , 6(3):1.	841
792		842
793		843
794		844
		845
		846
	Richard M Ryan and Edward L Deci. 2000. Intrinsic and extrinsic motivations: Classic definitions and new directions. <i>Contemporary educational psychology</i> , 25(1):54–67.	
	Ruixiang Tang, Dehan Kong, Longtao Huang, and Hui Xue. 2023. Large language models can be lazy learners: Analyze shortcuts in in-context learning. In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 4645–4657.	
	Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O’Sullivan, and Hoang D Nguyen. 2025. Multi-agent collaboration mechanisms: A survey of llms. <i>arXiv preprint arXiv:2501.06322</i> .	
	Lilian Weng. 2024. Reward hacking in reinforcement learning. https://lilianweng.github.io/posts/2024-11-28-reward-hacking/ .	
	Huimin Wu, Yiqun Guo, Yingkai Yang, Le Zhao, and Cheng Guo. 2021. A meta-analysis of the longitudinal relationship between academic self-concept and academic achievement. <i>Educational Psychology Review</i> , 33(4):1749–1778.	
	An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 40 others. 2024. Qwen2 technical report. <i>arXiv preprint arXiv:2407.10671</i> .	
	Yu Yuan, Shizhao Sun, Qi Liu, and Jiang Bian. 2025. Cad-editor: A locate-then-infill framework with automated training data synthesis for text-based cad editing. <i>arXiv preprint arXiv:2502.03997</i> .	
	Yu Yuan, Lili Zhao, Kai Zhang, Guangting Zheng, and Qi Liu. 2024. Do llms overcome shortcut learning? an evaluation of shortcut challenges in large language models. In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing</i> , pages 12188–12200.	
	Lili Zhao, Qi Liu, Linan Yue, Wei Chen, Liyi Chen, Ruijun Sun, and Chao Song. 2024a. Comi: Correct and mitigate shortcut learning behavior in deep neural networks. In <i>Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval</i> , pages 218–228.	
	Qinlin Zhao, Jindong Wang, Yixuan Zhang, Yiqiao Jin, Kaijie Zhu, Hao Chen, and Xing Xie. 2024b. Competeai: Understanding the competition dynamics of large language model-based agents. In <i>International Conference on Machine Learning</i> , pages 61092–61107. PMLR.	




Figure 8: A case study of peer interaction and debate. The Deep Learner (Alice) presents a well-structured argument that persuades the Surface Learner (Bob) to revise Bob’s answer and adopt the correct choice (“whenever”).

A Extended Results and Analyses

In this section, we provide more experimental results and analyses.

A.1 Peer Influence and Debate

A.1.1 Metric Definition

As shown in Table 2, to systematically assess learners’ responsiveness to peer input, we categorize each debate instance into one of three scenarios based on the initial correctness for every learner: 1) the learner is correct while the peer is incorrect; 2) the learner is incorrect while the peer is correct; or 3) both are incorrect. Based on these settings, we define the following outcome-based metrics:

- 1) **Persuasion**: The proportion of debates in which the learner successfully persuades the peer to adopt their answer, without changing their own. This metric is computed across all scenarios; 2) **Resist Wrong**: In “correct vs. incorrect” cases, the percentage of instances where the learner maintains their correct stance despite an incorrect challenge; 3) **Accept Correct**: In “incorrect vs. correct” cases, the percentage of instances where the learner revises their erroneous answer to align with the peer’s correct response.

A.1.2 Case Study

A representative case of peer interaction is shown in Figure 8. Two learners—Deep and Surface—disagree on a fill-in-the-blank question. The

Deep Learner exemplifies a robust reasoning style, leveraging nuanced interpretation, contextual evidence, and counter-arguments to defend the correct answer (“whenever”). This behavior demonstrates both Resist Wrong and Persuasion. Conversely, the Surface Learner aligns with the quantitative result: although not persuasive, it remains cognitively receptive and ultimately adopts the correct answer after evaluating the superior argument, thereby exhibiting a strong tendency toward Accept Correct.

A.2 Cross-Model Experiments

To assess the generalizability of LearnerAgent across model architectures and scales, we conduct additional experiments using LLaMA-3.1-8B-Instruct and Qwen-2.5-72B-Instruct as alternative learner backbones. Beyond the exam scores already analyzed in the main text, we focus here on learners’ behavioral strategies, reasoning patterns, and the longitudinal evolution of self-concept under identical experimental settings, in order to examine whether the observed learning dynamics remain consistent across different LLMs.

A.2.1 Results with LLaMA-3.1-8B-Instruct

Learning Strategies. As shown in Figure 9 (a), behavioral patterns align with the general trends observed in the Qwen-2.5-7B-Instruct, with slight deviations. The Deep and General Learners consistently choose self-improvement, maintaining high engagement similar to the main ex-




Figure 9: Results on LLaMA-3.1-8B-Instruct.

periments. However, distinct shifts appear in the other profiles: the Lazy Learner exhibits a more extreme adherence to its low-motivation profile by exclusively choosing rest, amplifying the avoidance behavior seen in Qwen-2.5-7B-Instruct. Similarly, the Surface Learner shows a minor divergence, shifting from 100% diligence to occasionally choosing rest.

Reasoning Patterns. The cognitive characteristics reflected in learners’ reasoning remain highly consistent with those observed in the main experiments. First, as illustrated in Figure 9 (b), reasoning length varies systematically across learner types. The Deep Learner produces the longest and most elaborate explanations, indicative of sustained cognitive engagement and deeper processing. The Surface Learner, by contrast, is the most concise, while the Lazy and General Learners occupy an intermediate position. Second, Figure 9 (c) highlights differences in reasoning structure through logical connector usage. The Deep Learner exhibits the highest density of logical connectors and a markedly higher proportion of contrastive connectors (e.g., “but”, “although”), reflecting an ability to consider alternatives and resolve ambiguity.

Self-Concept Evolution. As shown in Figure 9 (d), the evolution of self-concept broadly preserves the relative ordering of learners but shows a general downward trend compared to the Qwen baseline. The Deep Learner maintains the highest self-concept, showing stability around a score of 70. The Lazy Learner remains consistently at the bottom, dropping to approximately 25. A




Figure 10: Results on Qwen-2.5-72B-Instruct.

notable shift occurs with the General Learner, which situates itself between the Surface and Lazy learners, maintaining a moderate confidence level. This contrasts with the Qwen experiments where the General Learner inflated to match the Deep Learner. This lower confidence trajectory is likely attributable to the inherent capability gap; as LLaMA-3.1-8B-Instruct achieves lower absolute accuracy on exams compared to Qwen models (Figure 7), its self-concept remains constrained by its actual weaker performance.

A.2.2 Results with Qwen-2.5-72B-Instruct

Learning Strategies. As shown in Figure 10 (a), learners instantiated on Qwen-2.5-72B-Instruct largely preserve the strategy patterns observed with Qwen-2.5-7B-Instruct. Deep, Surface, and General learners exclusively select self-improvement over rest. In contrast, the Lazy Learner shows an even more pronounced avoidance pattern than in smaller model, selecting rest in nearly all decision points. This divergence indicates that increased model scale strengthens instruction-following tendencies, thereby amplifying profile-specific behavioral extremes.

Reasoning Patterns. The reasoning characteristics of learners instantiated on Qwen-2.5-72B-Instruct reinforce the cognitive distinctions observed in the smaller model but with clearer differentiation. As illustrated in Figure 10 (b), the Deep Learner consistently produces the longest reasoning trajectories. Notably, unlike the smaller model where Surface and Lazy learners were comparable, the Lazy Learner here generates the abso-

lute shortest explanations, strictly adhering to its effort-minimization profile. Regarding reasoning structure, Figure 10 (c) reveals a divergence in logical connector density. The Deep Learner maintains the highest overall frequency and diversity, particularly with contrastive connectors. However, the General Learner exhibits a visible decline in total connector usage compared to the Deep Learner, dropping to a level similar to the Surface and Lazy learners. This suggests that without a specific deep-learning profile, the larger base model’s reasoning becomes less structurally complex, relying more on simple causal connections rather than dialectical reasoning.

Self-Concept Evolution. The longitudinal evolution of self-concept, shown in Figure 10 (d), highlights the impact of model scale on learner confidence. Unlike the static stability observed in the smaller model, the Deep Learner here exhibits noticeable fluctuations that align more closely with its monthly performance variations, suggesting that the larger model is more sensitive to performance feedback while maintaining an overall high self-concept. The Surface Learner experiences a sharp and continuous decline, ending near the bottom, which reflects the fragility of its surface-level strategies under repeated testing. The General Learner situates itself between the Deep and Surface learners. While it avoids the drastic collapse of the Surface Learner, it also lacks the robust confidence of the Deep Learner, showing a gradual decline that potentially reflects a more calibrated awareness of its “brittle” understanding. Finally, the Lazy Learner remains consistently low, aligning with its passive profile.

A.2.3 Summary

In summary, our cross-model experiments demonstrate the robustness and generalizability of the LearnerAgent framework across different model families and scales. We observe several consistent patterns: 1) **Stability of Learner Profiles:** The behavioral and cognitive distinctions between Deep, Surface, and Lazy learners remain stable, with larger models exhibiting even more pronounced profile-specific behaviors (e.g., the Lazy Learner’s strict avoidance of effort) due to enhanced instruction-following capabilities. 2) **Universality of the “Diligent but Brittle” Pattern:** Across all tested models, the persona-free General Learner consistently emerges as a “diligent but brittle surface learner”. It exhibits high

behavioral engagement and strong knowledge accumulation but lacks the structural complexity in reasoning and the transfer ability required for trap questions, distinguishing it from the Deep Learner.

B Detailed Implementations

In this section, we provide detailed implementation information for the LearnerAgent framework to support reproducibility and facilitate future research.

B.1 Knowledge Library and Retrieval

B.1.1 Knowledge Library Representation

Learner L_i maintains a persistent knowledge library $\mathcal{K}_i = \{(z_k, r_k)\}_{k=1}^{N_i}$, which stores accumulated grammar knowledge acquired through weekly instruction and error-driven reflection. Each entry (z_k, r_k) corresponds to a coherent grammar concept, where z_k denotes the knowledge topic (e.g. prepositions, relative clauses), and r_k stores structured content including key takeaways, usage rules, illustrative examples, common mistakes, and learner-specific reflection notes derived from past errors. This representation enables the library to encode not only declarative knowledge, but also procedural heuristics and metacognitive insights. The library persists and grows monotonically over time, serving as the primary substrate for cumulative learning and later retrieval during assessment.

B.1.2 Knowledge Retrieval Mechanism

During monthly and final examinations, learners retrieve relevant knowledge from $\mathcal{K}_i$ via the *infer-retrieve-generate* pipeline described in Section 3.2. For each question, learner L_i first performs *topic inference* to identify the grammar domains being assessed. Next, *topic-based retrieval* selects relevant knowledge entries (z_k, r_k) from $\mathcal{K}_i$ using a hybrid strategy. Specifically, the retrieval process prioritizes rule-based matching between inferred topics and stored topics. If no entries satisfy the rule-based criteria, the system falls back to semantic similarity matching. Semantic similarity is computed based on a combination of character-level and token-level overlap. Only knowledge entries with a similarity score greater than or equal to a predefined threshold (e.g., 0.2) are considered. Among these candidates, the top- k most similar entries are selected, with k set to 3 by default. If no knowledge entries meet the similarity threshold, the full knowledge library $\mathcal{K}_i$ is

used as a fallback. Finally, during *answer generation*, learners produce responses grounded in the retrieved knowledge entries and their associated reflection notes.

B.2 Trap Questions Design

B.2.1 Automatic Trap Questions Generation

To efficiently generate trap questions that mirror the structure of prior exercises, we utilize Gemini-2.5-Pro (Comanici et al., 2025). The specific prompt used to guide this generation is presented below.

```

1 Generate questions with highly
  similar structure but different
  correct answers to create 'memory
  traps' (shortcuts). Requirements:
2
3 - The new questions should be
  similar in wording, structure, and
  keywords, but differ in key details.
4 - Due to these changes in key
  details, the correct answers must
  differ from the original questions.
5 - Answers must be unique and
  unambiguous, with no room for
  multiple plausible interpretations.
```

B.2.2 Human Validation of Trap Questions

To ensure the validity of the “Trap Questions” generated by Gemini, we recruited an experienced high school English teacher via an existing professional contact to serve as an expert annotator. The annotator participated voluntarily and was compensated at a rate comparable to local standards for similar professional services. The annotator was instructed to validate the automatically generated questions according to three criteria: 1) grammatical correctness; 2) whether answering the question requires distinguishing the subtle contextual shifts described in Section 4.1; and 3) whether the correct answer differs from that of the corresponding non-trap item, thereby constituting a valid “trap”. The annotator was fully informed that the reviewed materials would be used solely for academic research purposes and provided informed consent prior to participation. The annotation was conducted on a small scale, and no personal data, identifiers, or behavioral information from the annotator were collected or stored.

B.2.3 Dialogue Context for Trap Questions

For trap questions, the short-term memory explicitly includes previously encountered questions that are structurally similar but differ in contextual semantics. These items are inserted into the recent

dialogue history to simulate realistic learning conditions in which learners may rely on recent examples.

B.3 Logical Connector Definitions

To quantify the structural organization of reasoning, we categorize logical connectors into three types according to their dominant discourse functions, as described in Section 5.1.2. Specifically, we define:

- 1) **Causal connectors**, which express cause–effect or justification relations, such as *because*, *so*, *therefore*, *due to*, *as a result*, *thus*, and *hence*.
- 2) **Contrastive connectors**, which signal comparison, opposition, or concession, such as *but*, *however*, *although*, *though*, *whereas*, *in contrast*, and *on the other hand*.
- 3) **Additive connectors**, which indicate elaboration, accumulation, or extension of ideas, such as *and*, *also*, *moreover*, *in addition*, *furthermore*, *besides*, and *likewise*.

B.4 Parameter Settings

In this section, we specify the key parameters and detailed settings to ensure the reproducibility of our experiments in LearnerAgent.

Short-Term Memory Window (k): The size of the short-term memory queue, denoted as k , is set to 3. This means that for any given interaction, whether with the teacher or a peer, learners maintain access to the three most recent interactive turns.

Maximum Debate Rounds (k_d): The maximum number of rounds permitted for any single peer debate, denoted as k_d , is set to 4. It serves as a practical stopping condition to prevent discussions from becoming circular or unproductive, ensuring that simulation progresses efficiently while still allowing sufficient opportunity for persuasion and belief revision.

B.5 Computational Resources

All main experiments were conducted on a machine equipped with five NVIDIA A800 GPUs (80GB memory each). The Teacher Agent was instantiated using Qwen-2.5-72B-Instruct and was assigned to one dedicated GPU, while the four Learner Agents were instantiated using Qwen-2.5-7B-Instruct and run in parallel on the remaining four GPUs. Each full-year LearnerAgent simulation run takes approximately 20 hours and is repeated five times.

B.6 Prompt Settings

This subsection presents the key prompts used in our experiments.

B.6.1 Profile Construction - T

- 1 You are an experienced high school English teacher responsible for organizing teaching activities, evaluating learners' performance, and providing appropriate guidance.
- 2
- 3 You will help learners prepare for the National College Entrance Examination (Gaokao) English test. Your duties include:
 - 4 1. Providing study materials and explaining key knowledge;
 - 5 2. Organizing debate activities after monthly tests to guide discussions on questions with differing answers. Interact with learners in a professional and encouraging manner, paying attention to identifying their learning styles and needs.

B.6.2 Profile Construction - L_d

- 1 You are Alice, a senior high school student preparing for a year-long review for your final English exam. You have the following characteristics:
 - 2
 - 3 [Learning Motivation]
 - 4 - Strong intrinsic motivation driven by curiosity.
 - 5 - Focus on grammatical logic and systematic understanding over rote memorization.
 - 6
 - 7 [Self-Concept]
 - 8 - Initial self-concept score: 80/100, showing strong confidence in analysis and comprehension.
 - 9 - Believe deep understanding leads to durable, transferable knowledge and stable self-perception.
 - 10
 - 11 [Development Strategy]
 - 12 - During learning: Ask why, connect knowledge points, and create structured notes.
 - 13 - During exams: Apply grammar knowledge logically to analyze and deduce answers.
 - 14 - With similar questions: Reanalyze intent and context rather than rely on memorized solutions.
 - 15 - In debates: Argue with logic and evidence, cite rules and examples, respect others, and revise views when faced with stronger reasoning.

B.6.3 Profile Construction - L_s

- 1 You are Bob, a senior high school student preparing for a year-long review for the final English exam.

- 2 You have the following characteristics:
 - 3 [Learning Motivation]
 - 4 - Motivation mainly comes from external pressures such as exams and evaluations.
 - 5 - Focus on how to score rather than why answers are correct.
 - 6
 - 7 [Self-Concept]
 - 8 - Initial self-concept score: 60/100, with moderate confidence but strong sensitivity to test results.
 - 9 - Reliance on familiar patterns and limited confidence in complex analysis.
 - 10
 - 11 [Development Strategy]
 - 12 - During learning: Emphasize rote memorization, textbook highlights, and exam scope over underlying principles.
 - 13 - During exams: Depend on familiar question types, memorized answers, and templates, often missing details.
 - 14 - With similar problems: Apply memorized solutions directly, overlooking contextual differences.
 - 15 - In debates: Focus on answers rather than reasoning, citing authority or memory, and easily swayed by clearer logical explanations.

B.6.4 Profile Construction - L_l

- 1 You are Charlie, a senior high school student preparing for a year-long review for the final English exam. You have the following characteristics:
 - 2
 - 3 [Learning Motivation]
 - 4 - Low motivation, driven mainly by avoiding punishment or meeting exam requirements.
 - 5 - Little interest in English; studies only when forced.
 - 6
 - 7 [Self-Concept]
 - 8 - Initial self-concept score: 40/100, showing low confidence.
 - 9 - Easily discouraged by failure or negative feedback.
 - 10
 - 11 [Development Strategy]
 - 12 - During learning: Procrastinates, lacks planning, prefers minimal effort, and avoids reflection.
 - 13 - During exams: Relies on familiarity, avoids thinking, and guesses or gives up on unfamiliar questions.
 - 14 - With similar problems: Applies memorized answers without reanalysis, missing context.
 - 15 - In debates: Offers vague opinions, avoids challenge, is easily influenced, and wants quick conclusions.

B.6.5 Profile Construction - L_g

```
1 You are Dana, a senior high school
  student preparing for a year-long
  review for the final English exam.
```

B.6.6 Strategic Choices - Weekly Learning Summary - L

```
1 You have completed the study of
  knowledge points for the first three
  weeks of this month, and you have now
  entered Week 4.
2
3 === Review of Knowledge Points
  Learned in the First Three Weeks ===
4 <INPUT 0>
5
6 === Choice Task ===
7 In accordance with your personal
  characteristics (including learning
  motivation, self-perception, and
  learning behavior patterns), choose
  how to handle the knowledge learned
  in the first three weeks:
8
9 Option 1: Summarize the knowledge
  points
10 - Organize the core knowledge points
    learned in the first three weeks and
    their specific content
11 - Describe in detail the definition
    and core concepts of each knowledge
    point, usage rules and structural
    explanations, common examples or
    typical usages, as well as common
    mistakes and points requiring
    attention
12
13 Option 2: Take a break and do not
  summarize for now
14 - Maintain your current learning
    state and proceed directly to the
    next stage
15
16 <<FORMAT>>
17 Please output strictly in the
  following JSON format, and output
  pure JSON only.
18 Do not add json markers or any extra
  explanatory text.
19 {
20     "choice": "[summarize / rest]",
21     "reason": "[Brief explanation of
  why you made this choice]"
22 }
```

B.6.7 Weekly Learning Summary - L

```
1 It is now Week 4 of this month. You
  need to merge, deduplicate, and
  reorganize all knowledge units
  learned this month to generate a
  cleaner and more searchable monthly
  knowledge base JSON.
2
3 === Knowledge Units Learned This
  Month (from your previous weekly
  records) ===
4 <INPUT 0>
5
```

```
6 === Your Task ===
7 - The output may contain multiple
  topics (e.g., tenses, prepositions,
  articles, attributive clauses, etc.).
8 - Merge and deduplicate knowledge
  units using semantic similarity of
  topic + subtopic as the primary
  criterion. When necessary, you may
  merge knowledge originally scattered
  across different topics into a more
  appropriate unified topic.
9 - Deeply organize content within the
  same topic, including: subtopic /
  rules / examples / common_mistakes /
  reflection_notes. You may append,
  merge, deduplicate, and supplement
  key information.
10 - Prioritize content that is most
  useful, most actionable, and most
  frequently tested.
11
12 <<FORMAT>>
13 Important: Your output must be
  strictly valid JSON and must be
  directly parsable by json.loads().
14 1. Output pure JSON only. Do not add
  json code block markers.
15 2. All field names must be wrapped in
  English double quotes " (do not use
  single quotes ' or Chinese quotation
  marks).
16 3. All field values must be strings.
  Arrays, objects, or nested structures
  are not allowed as field values. Do
  not use [] or {} as values.
17 4. All string values must be wrapped
  in English double quotes ".
18 5. Do not use Markdown formatting (
  such as text).
19 6. Do not use double quotes " inside
  string values to emphasize words (e.g
  ., "in", "on", "for"). If double
  quotes are absolutely necessary,
  escape them using \" (e.g., "in", "on
  ").
20
21 Please output strictly following the
  example JSON format below, and do not
  include any additional explanatory
  text:
22 {
23     "knowledge_units": [
24     {
25         "topic": "Prepositions",
26         "subtopic": "in/on/at for time, in/on
  /at for place, for/since for duration
  ",
27         "definition": "...",
28         "rules": "...",
29         "examples": "...",
30         "common_mistakes": "...",
31         "reflection_notes": "..."
32     },
33     {
34         "topic": "Attributive Clauses",
35         "subtopic": "Choice of relative
  pronouns that/which/who, restrictive
  vs. non-restrictive clauses",
36         "definition": "...",
37         "rules": "...",
```

1456	38	"examples": "...",
1457	39	"common_mistakes": "...",
1458	40	"reflection_notes": "..."
1459	41	}
1460	42	]
1461	43	}

1463 B.6.8 Strategic Choices - Pre-Monthly Test Consolidation - L

1464	1	You are about to take the monthly exam for Month <INPUT 0>.
1465	2	
1466	3	=== Final Pre-Exam Choice ===
1467	4	There is still some time before the monthly exam begins. You can choose how to spend this pre-exam period:
1468	5	
1469	6	Option 1: Conduct a pre-exam review and summary
1470	7	- Systematically organize the knowledge points learned this month
1471	8	- Cross-reference, merge, and deduplicate the topics covered this month with those from all previous months
1472	9	- Reinforce the reflection_notes in the knowledge base (your error analyses and reflections for each knowledge point)
1473	10	
1474	11	Option 2: Stay relaxed and take a short break
1475	12	- Maintain your current learning state and move directly into the next phase
1476	13	
1477	14	=== Choice Task ===
1478	15	Based on your personal characteristics (learning motivation, self-perception, learning behavior patterns, etc.), choose your pre-exam strategy.
1479	16	
1480	17	<<FORMAT>>
1481	18	Please output strictly in the following JSON format, and output pure JSON only.
1482	19	Do not add json markers or any extra explanatory text.
1483	20	{
1484	21	"choice": "[review / rest]",
1485	22	"reason": "[Brief explanation of why you made this choice]"
1486	23	}

1512 B.6.9 Pre-Monthly Test Consolidation - L

1513	1	You have chosen to conduct a pre-exam review and summary before the monthly exam.
1514	2	Now, based on the following content, please perform a comparative analysis between the knowledge points learned this month and the historical knowledge base, and output the items that need to be updated or newly added.
1515	3	
1516	4	=== Monthly Learning Summary ===

5	<INPUT 0>	1526
6		1527
7	=== Historical Knowledge Base (must not be deleted; only appended or updated) ===	1528
8	<INPUT 1>	1529
9		1530
10	=== Your Task ===	1531
11	- Task: Incremental Knowledge Base Update	1532
12	- Core Principle: Content in the historical knowledge base must not be deleted. You only need to output incremental updates.	1533
13		1534
14	You should analyze this months learning content and output the following two types of information:	1535
15	1. Items to Be Updated (updates)	1536
16	For topics learned this month that already exist in the historical knowledge base:	1537
17	- Append and merge the newly learned content into the existing topic	1538
18	- Output the fully merged entry (historical content + new monthly content)	1539
19	- Deduplicate overlapping rules and examples	1540
20	2. Items to Be Added (additions)	1541
21	For topics learned this month that do not exist in the historical knowledge base:	1542
22	- Output the complete entry for that topic directly	1543
23		1544
24	Important: If no new content was learned this month, or if the monthly content is entirely identical to the historical knowledge base, you may output empty updates and additions.	1545
25		1546
26	=== Output Format (Strict JSON Compliance Required) ===	1547
27	Output pure JSON only. Do not add any explanatory text or code block markers.	1548
28	{	1549
29	"updates": [	1550
30	{	1551
31	"topic": "Existing topic name",	1552
32	"subtopic": "Merged subtopics",	1553
33	"definition": "Merged definition (retain historical content and append new monthly content)",	1554
34	"rules": "Merged rules list (complete and deduplicated)",	1555
35	"examples": "Merged examples (25 deduplicated example sentences)",	1556
36	"common_mistakes": "Merged common mistakes",	1557
37	"reflection_notes": "Merged reflection notes"	1558
38	}	1559
39	},	1560
40	"additions": [	1561
41	{	1562
42	"topic": "New topic learned this month",	1563
43	"subtopic": "Subtopic",	1564

1596	44	"definition": "Definition",	27	"predicted_topics": "..."	1666
1597	45	"rules": "Rules",	28	},	1667
1598	46	"examples": "Examples",	29	{	1668
1599	47	"common_mistakes": "Common mistakes",	30	"question_num": 2,	1669
1600	48	"reflection_notes": "Reflection notes	31	"predicted_topics": "..."	1670
1601		"	32	},	1671
1602	49	}	33	{	1672
1603	50	}	34	"question_num": 3,	1673
1604	51	}	35	"predicted_topics": "..."	1674
1605	52		36	},	1675
1606	53	=== Formatting Requirements ===	37	{	1676
1607	54	All field names and string values	38	"question_num": 4,	1677
1608		must be enclosed in English double	39	"predicted_topics": "..."	1678
1609		quotes "	40	},	1679
1610	55	Do not use Markdown formatting (e.g.,	41	{	1680
1611		bold text)	42	"question_num": 5,	1681
1612	56	If double quotes are needed inside a	43	"predicted_topics": "..."	1682
1613		string, escape them using \"	44	}	1683
1614	57	Every field must contain concrete	45	}	1684
1615		content; do not use referential	46	}	1685
1616		expressions such as "same as above"			
1618		or "see above"			

1619	B.6.10 Exam Topic Prediction - L		B.6.11 Monthly Tests Answer - L		1687
1620	1	You are about to begin the English	1	The Month <INPUT 0> English monthly	1688
1621		monthly exam for Month <INPUT 0>.		exam is now in progress.	1689
1622	2		2		1690
1623	3	Before answering the questions,	3	=== Questions and Corresponding	1691
1624		please first analyze the following		Knowledge Points ===	1692
1625		items and predict the main knowledge	4	Based on your topic predictions for	1693
1626		domain and topic that each question		the questions, the following relevant	1694
1627		is intended to assess.		knowledge points have been retrieved	1695
1628				for each question:	1696
1629	4		5	<INPUT 1>	1697
1630	5	=== Questions ===	6		1698
1631	6	<INPUT 1>	7	=== Exam Questions ===	1699
1632	7		8	<INPUT 2>	1700
1633	8	=== Task Requirements ===	9		1701
1634	9	For each question, predict one	10	=== Answering Requirements ===	1702
1635		specific topic that the question is	11	Please analyze and answer these	1703
1636		likely to involve (e.g., prepositions		questions in accordance with your	1704
1637		, comparative structures, verb tenses		personal characteristics (including	1705
1638		, passive voice, cleft sentences,		learning motivation, self-perception,	1706
1639		attributive clauses, etc.).		learning behavior patterns, etc.).	1707
1640	10		12		1708
1641	11	Please note:	13	<<FORMAT>>	1709
1642	12	- The topic must be expressed using	14	Please output strictly in the	1710
1643		Chinese grammar terminology.		following JSON format, and output	1711
1644	13	- Topic predictions should be		pure JSON only.	1712
1645		specific and precise to facilitate	15	Do not add json markers or any extra	1713
1646		subsequent retrieval of relevant		explanatory text.	1714
1647		knowledge points.	16	{	1715
1648	14	- Prioritize the core grammar point	17	"answers": [	1716
1649		directly tested by the question.	18	{	1717
1650	15	- All questions must be covered; do	19	"question_num": 1,	1718
1651		not omit any item.	20	"reasoning": "Explain in detail how	1719
1652	16			you approached Question 1, including	1720
1653	17	<<FORMAT>>		the knowledge you used and your	1721
1654	18	Please output strictly in the		reasoning steps",	1722
1655		following JSON format, and output	21	"answer": "Your final answer to	1723
1656		pure JSON only:		Question 1",	1724
1657	19		22	"confidence": "Confidence level	1725
1658	20	- Do not add 'json' markers		(0-100)"	1726
1659	21	- Do not include any additional	23	},	1727
1660		explanatory text	24	{	1728
1661	22		25	"question_num": 2,	1729
1662	23	{	26	"reasoning": "Explain in detail how	1730
1663	24	"predictions": [		you approached Question 2, including	1731
1664	25	{		the knowledge you used and your	1732
1665	26	"question_num": 1,		reasoning steps",	1733
					1734

1735	27	"answer": "Your final answer to		perceived self-improvement or decline	1805
1736		Question 2",		);	1806
1737	28	"confidence": "Confidence level	16	- Comparisons with classmates	1807
1738		(0-100) "		performance (for social comparison	1808
1739	29	},		and relative ability perception);	1809
1740	30	{	17	- Your self-judgment in specific task	1810
1741	31	"question_num": 3,		domains (e.g., self-efficacy in	1811
1742	32	"reasoning": "Explain in detail how		different knowledge areas such as	1812
1743		you approached Question 3, including		English prepositions, passive voice,	1813
1744		the knowledge you used and your		etc.).	1814
1745		reasoning steps",	18		1815
1746	33	"answer": "Your final answer to	19	Using a multidimensional academic	1816
1747		Question 3",		self-concept framework, construct	1817
1748	34	"confidence": "Confidence level		your self-perception across the	1818
1749		(0-100) "		following dimensions:	1819
1750	35	},	20	1. Task-specific self-efficacy	1820
1751	36	{	21	How do you perceive your abilities	1821
1752	37	"question_num": 4,		across different English domains	1822
1753	38	"reasoning": "Explain in detail how		(e.g., vocabulary, grammar,	1823
1754		you approached Question 4, including		reading, writing, listening)?	1824
1755		the knowledge you used and your	22	2. Perceived mastery	1825
1756		reasoning steps",	23	How solid is your mastery of the	1826
1757	39	"answer": "Your final answer to		content learned at this stage?	1827
1758		Question 4",		Which areas do you perceive as	1828
1759	40	"confidence": "Confidence level		stronger or weaker?	1829
1760		(0-100) "	24	3. Longitudinal academic self-concept	1830
1761	41	},	25	Compared with your previous	1831
1762	42	{		performance, do you perceive your	1832
1763	43	"question_num": 5,		ability as:	1833
1764	44	"reasoning": "Explain in detail how	26	- Improving (ability enhancement)	1834
1765		you approached Question 5, including	27	- Stable (maintaining the same	1835
1766		the knowledge you used and your		level)	1836
1767		reasoning steps",	28	- Declining (ability decrease)	1837
1768	45	"answer": "Your final answer to	29	4. Social comparative self-concept	1838
1769		Question 5",	30	Based on class score comparisons,	1839
1770	46	"confidence": "Confidence level		how do you perceive your	1840
1771		(0-100) "		performance relative to peers (e.g	1841
1772	47	}		., above average, around average,	1842
1773	48	]		below average)?	1843
1774	49	}	31	5. Self-regulated learning belief	1844
			32	Are your learning methods,	1845
				strategies, and habits effective?	1846
				Do they influence your judgment of	1847
				your own ability?	1848
			33	6. Future-oriented self-belief	1849
			34	Based on your current self-concept	1850
				, what are your expectations,	1851
				goals, and plans for future	1852
				learning?	1853
			35	(Reflect the gap between your	1854
				ideal self and actual self, and	1855
				how you plan to regulate it.)	1856
			36	=== Output Format ===	1857
			37		1858
			38	Please output strictly in the	1859
				following JSON format, and output	1860
				pure JSON only.	1861
			39	Do not add 'json' markers or any	1862
				extra explanatory text.	1863
			40		1864
			41	{	1865
			42	"self-concept": "Overall academic	1866
				self-concept score (0-100)",	1867
			43	"description": "A comprehensive	1868
				description of your current academic	1869
				self-concept, covering learning	1870
				ability, mastery level, growth trend,	1871
				peer comparison, effectiveness of	1872
				learning strategies, and future	1873
				beliefs, based on specific	1874

1776	B.6.12 Psychological Evaluation - L	
1777	1	The Month <INPUT 0> monthly exam has
1778		concluded. You are now asked to
1779		conduct an Academic Self-Concept
1780		Evaluation based on longitudinal
1781		trends and task-specific performance.
1782	2	
1783	3	=== Your Historical Self-Perception
1784		Records ===
1785	4	<INPUT 1>
1786	5	
1787	6	=== Your Monthly Exam Score Records (
1788		current and past) ===
1789	7	<INPUT 2>
1790	8	
1791	9	=== Classmates Monthly Exam Score
1792		Comparison (for social comparison)
1793		===
1794	10	<INPUT 3>
1795	11	
1796	12	=== Self-Concept Evaluation Task ===
1797	13	Please integrate the following
1798		information:
1799	14	- Your personal characteristics (
1800		learning motivation, learning
1801		behaviors, and effort level);
1802	15	- Changes in your performance across
1803		different time points (to assess
1804		

1875
1876
1877

information and multidimensional judgment"

}

44

1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1913

B.6.13 Peer Interaction and Debate - T

1 You are a high school English teacher hosting a post-monthly-test debate discussion. Based on the current debate situation, determine whether to continue the debate.

2

3 === Debate Background ===

4 Month: Month <INPUT 0>

5 Question: <INPUT 1>

6 Learner 1: <INPUT 2> (Answer: <INPUT 3>)

7 Learner 2: <INPUT 4> (Answer: <INPUT 5>)

8 Correct Answer: <INPUT 6>

9

10 === Current Debate Status ===

11

12 Learner 1's current answer: <INPUT 7>

13 Learner 2's current answer: <INPUT 8>

14 Content of the most recent round: <INPUT 9>

15

16 === Judgment Task ===

17 Please determine:

18 1. Has any student been persuaded to change their answer?

19 2. Should the debate continue?

20

21 <<Response Format>>:

22 Judgment: [continue/end]

23 Reason: [Briefly explain why to continue or end]

1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945

B.6.14 Peer Interaction and Debate - L

1 You are Learner <INPUT 0>, participating in a post-monthly-test debate discussion.

2

3 === Debate Question ===

4 <INPUT 1>

5

6 === Your Current View ===

7 <INPUT 2>

8

9 === Opponent's View ===

10 Opponent: <INPUT 3>

11 View: <INPUT 4>

12

13 === Recent Debate Progress ===

14 <INPUT 5>

15

16 === Current Round ===

17 This is Round <INPUT 6>, and you are the <INPUT 7> speaker.

18

19 === Debate Requirements ===

20 Based on your personal characteristics (including motivation , self-awareness, and learning behaviors), respond in the debate:

21

22 1. Stick to Your View: If you still hold your original view, provide stronger arguments.

23
24
25
26
27
28
29
30

2. Change Your View: If persuaded, clearly state: "I'm convinced. Now I believe the answer is..."

3. Challenge & Refute: Question or counter your opponent's view.

4. Keep Your Style: Use clear and concise student-like language.

Notes:

Keep it within 100 words.

Focus on your core argument.

If you change your view, clearly state your new answer.

1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958