

Agent4Edu: Generating Learner Response Data by Generative Agents for Intelligent Education Systems

Weibo Gao¹, Qi Liu^{1,2*}, Linan Yue¹, Fangzhou Yao¹, Rui Lv¹, Zheng Zhang¹, Hao Wang¹, Zhenya Huang^{1,2}

¹ State Key Laboratory of Cognitive Intelligence, University of Science and Technology of China

² Institute of Artificial Intelligence, Hefei Comprehensive National Science Center
 {weibogao, lnyue, fangzhouyao, lvru2018, zhangzheng}@mail.ustc.edu.cn
 {qiliuql, wanghao3, huangzhy}@ustc.edu.cn

Abstract

Personalized learning represents a promising educational strategy within intelligent educational systems, aiming to enhance learners' practice efficiency. However, the scarcity of offline practice response data (e.g., answer correctness) and potential biases in human online practice create a significant gap between offline metrics and the actual online performance of personalized learning services. To address this challenge, we introduce Agent4Edu, a novel personalized learning simulator leveraging recent advancements in human intelligence through large language models (LLMs). Agent4Edu features LLM-powered generative agents equipped with learner profile, memory, and action modules tailored to personalized learning algorithms. The learner profiles are initialized using real-world response data, capturing practice styles and cognitive factors. Inspired by psychology theory, the memory module records practice facts and high-level summaries, integrating reflection mechanisms. The action module supports various behaviors, including exercise understanding, analysis, and response generation. Each agent can interact with personalized learning algorithms, such as computerized adaptive testing, enabling a multifaceted evaluation and enhancement of customized services. Through a comprehensive assessment, we explore the strengths and weaknesses of Agent4Edu, emphasizing the consistency and discrepancies in responses between agents and human learners.

1 Introduction

Intelligent education platforms like *Coursera.com* and *LeetCode.com* provide a rich array of learning resources, such as courses and exercises, within a flexible online environment. The accessibility and convenience of these platforms have attracted a growing number of learners. A key online learning activity is "practice", where learners independently select and answer exercises. The platforms record their responses, such as the correctness of their answers. By analyzing response data, many personalized learning services, such as exercise recommendations, knowledge tracing, and computerized adaptive testing, can be tailored to meet each learner's specific needs, enhancing the learning process and increasing learner satisfaction. For instance, on *LeetCode*,

analyzing a learner's historical programming experiences allows the platform to recommend exercises of appropriate difficulty levels, thus optimizing learning gains.

The effectiveness of personalized learning services hinges on the availability of high-quality response data for the corresponding algorithm training. However, the scarcity of offline response data and potential biases in its correlation with online practice introduces a significant gap between offline metrics and actual online performance. This discrepancy impedes the integration of research with real-world applications. To bridge this gap, a promising approach is to simulate learner response data. Imagine an online platform equipped with a configurable simulation system that faithfully captures human learners' response patterns while seamlessly interacting with personalized learning algorithms. Such a simulator undoubtedly has the potential to revolutionize the traditional research paradigm in intelligent education, providing innovative avenues for response data collection, personalized algorithm development and evaluation.

Several approaches to simulating learner response data have been proposed and have achieved notable success (Piech et al. 2015; Zhao et al. 2023). However, two major limitations exist in current approaches: (1) **Simplified Simulations**. Most existing studies predict learners' responses (e.g., correct or incorrect answers) without considering the detailed answer processes by which humans use their knowledge to understand, analyze, and solve problems. Hence, these simulations may lack reliability and interpretability. (2) **Dependency on Real Response Data**. An ideal simulator should be capable of simulating learner responses even when real-world datasets are insufficiently available, thereby enhancing its applicability. However, current methods require high-quality real-world data to train the simulation strategy. As a result, these methods can only generate learner response data similar to existing real-world datasets and struggle to generalize to more challenging scenarios, such as zero-shot simulations.

Recent advancements in large language models (LLMs) have demonstrated remarkable capabilities in autonomous interaction and decision-making (Brown et al. 2020; Ouyang et al. 2022; Yue et al. 2023; Jin et al. 2023; Long et al. 2024). These advancements underscore the potential of leveraging LLM-powered agents to simulate human social behav-

*Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

iors, such as daily life in Smallville (Park et al. 2023) and software development (Qian et al. 2023). LLM-based user simulators possess rich pre-trained knowledge and human-like intelligence, enabling them to perceive and simulate intricate human practice processes. Furthermore, their in-context learning ability allows LLMs to perform zero-shot simulations with minimal reliance on real-world data (Wang et al. 2023c). Consequently, LLM-based generative agents present a promising approach for addressing the current limitations of learner response simulators.

In this paper, we introduce **Agent4Edu**, a personalized learning simulator designed for intelligent educational systems, comprising two key components: an **LLM-powered generative agent** and a **personalized learning environment** (see the framework in Figure 1). From a learner perspective, the LLM-powered generative agent is responsible for simulating learners’ response data by capturing their response patterns and inferring problem-solving actions. Each agent is initialized based on available learner response data and consists of three modules: a learner profile, memory, and action module. The learner profile module stores learners’ past practice styles (e.g., *activity*) and cognitive factors (e.g., *ability*), aligning with human learners’ learning status. The memory module, inspired by psychological theories (Baker 2001) and human learning mechanism (Wang et al. 2023d), records past practice experiences and summarizes learning status through reflections. This facilitates coherent observations, monitors *knowledge proficiency evolution*, reinforces memory, and simulates human forgetting. The action module enables agents to choose, understand, analyze, and solve exercises recommended by personalized learning algorithms, leading to more reliable and interpretable response generation. Our agent can also utilize tools, such as employing the psychological IRT model (Baker 2001) to assess *ability* within the Profile module and using DNeural-CDM (Wang et al. 2023a) to trace *knowledge proficiency evolution* within the Memory module. From a personalized learning perspective, the learning environment can be configured with any personalized learning algorithm, allowing agents to interact directly and simulate a real learning environment. Notably, despite extensive research on simulating user behavior with generative agents, we are the first to focus specifically on educational scenarios to generate response data for individual learners.

Our main contributions are summarized as follows:

- We develop Agent4Edu¹, a personalized learning simulator that leverages LLM-powered generative agents to simulate human learners’ response data as well as demonstrate the practice process. Additionally, the agent interacts with personalized learning environments to evaluate and improve intelligent tutoring algorithms.
- Our generative agents, featuring profile, memory, and action modules specifically designed for “Education”, can not only generate response data but also accurately simulate human choices, understanding, analysis, and

problem-solving for exercises, outperforming existing learner simulation methods.

- To systematically evaluate Agent4Edu, we conduct comprehensive experiments from both the agent and personalized learning perspectives. From the agent perspective, we assess the consistency between the agents and human learners. From the learning perspective, we evaluate and improve personalized learning algorithms for computerized adaptive testing, based on generative agents and simulated data. Extensive experimental results demonstrate the effectiveness of Agent4Edu.

2 Related Work

Learner Response Data Simulation Learner Simulation aims to address the shortage of high-quality practice data in intelligent educational systems and has been applied in numerous previous studies (Zhao et al. 2023; Yao et al. 2024). Memory-based (Reddy, Levine, and Dragan 2017) relies on manually crafted rules to predict learners’ responses or memory behavior. EERNN (Su et al. 2018) and KES (Liu et al. 2019) utilize RNN-based models to forecast learners’ performance. DAISim (Zhao et al. 2023) constructs learner simulations as Markov decision processes, simultaneously considering learners’ long and short-term question-answering patterns. However, the memory-based simulator is overly simplistic and cannot simulate complex interactions. Most other learner simulators simplify the student answering process and face challenges in conducting zero-shot simulations due to their reliance on data. In this paper, we employ an LLM-powered agent to simulate the student practice process, addressing these limitations.

Personalized Learning Services Intelligent educational systems offer learners personalized learning services, including Computerized Adaptive Testing (CAT) (Chang and Ying 1996), exercise recommendation (Huang et al. 2019) and learning path suggestions (Liu et al. 2019), to help learners enhance their skills. In this work, we select the representative and popular CAT services as our personalized learning scenarios for our study and experiments. CAT is an advanced educational measurement method that evaluates the knowledge level of examinees in minor exercises, which has been widely used in various standardized tests (e.g., GMAT and GRE) (Zhuang et al. 2024; Bi et al. 2020; Lord 2012; Chang and Ying 1996). However, current CAT models require high-quality practice data to train a cognitive diagnosis model for evaluating learner ability or knowledge proficiency, which is often challenging to gather. Therefore, in this paper, we employ the CAT service within learning systems to assess the quality of the data generated by our agents. Additionally, we investigate the potential for enhancing CAT models using simulated data.

LLM-based Agents LLM-based generative agents demonstrate the remarkable capabilities to perceive their environment, make decisions, and take actions, thus, emerging a substantial amount of research (Wang et al. 2024b). The development of generative agents (Park et al. 2023), designed with profile, memory, action, and reflective capabilities, represents pioneering work in simulating human

¹The code, data, and appendix are publicly available at <https://github.com/bigdata-ustc/Agent4Edu>.

daily life. Within this general framework, agents tailored to specific tasks (Qian et al. 2023; Wu et al. 2023; Wang et al. 2023b; Huang et al. 2023; Zhang et al. 2023b,a) and simulations (Gao et al. 2023; Wang et al. 2023d; Park et al. 2023; Liu et al. 2023; Wang et al. 2023c) have been constructed. Recent research highlights bringing generative agents to educational settings (Li et al. 2024; Dan et al. 2023; Kieser et al. 2023). For example, (Qadir 2023; Rahman and Watanobe 2023) conclude the applications of ChatGPT to engineering education. (Baidoo-Anu and Ansah 2023) focus on the literature review over the published paper. SocraticLM (Liu et al. 2024) embodies a “Thought-Provoking” teaching paradigm, engaging students in active problem-solving, akin to a real classroom teacher. The most relevant part of our work is EduAgent (Xu, Zhang, and Qin 2024) which utilizes LLM-based agents to simulate learners studying PowerPoint presentations and videos, predicting their quiz outcomes to assess performance. However, this approach relies on expert-annotated cognitive factors to initialize agents, disregarding the understanding and analysis of exercises. In contrast, our Agent4Edu extracts cognitive factors from data using tools and captures practice styles, allowing it to simulate the detailed exercise understanding and analysis process and interact effectively with personalized learning algorithms.

3 Agent4Edu

Agent4Edu is a personalized learning simulator, aimed at accurately simulating learners’ response data and facilitating responsive personalized learning algorithms. It contains two key components: (1) LLM-powered generative agents that capture learners’ practice patterns and cognitive preferences to simulate their response, and (2) a personalized learning environment that interacts with agents to support accurate and interpretable evaluations and improvements of mainstream intelligent algorithms (e.g., computerized adaptive testing). The framework of Agent4Edu is illustrated in Figure 1. All the prompts are listed in Appendix C.

3.1 Task Formulation

Suppose there are $|U|$ learners, $|E|$ exercises in an intelligent educational system. For a learner $u \in U$, his/her response data are denoted as a time-ordered set $l_u = \{(e_1, c_{e_1}, y_{u,e_1}), (e_2, c_{e_2}, y_{u,e_2}), \dots, (e_n, c_{e_n}, y_{u,e_n})\}$, where $e_i \in E$ represents the exercise that learner u practiced at step i , and y_{u,e_i} is u ’s response to exercise e_i , which is usually denoted as a binary value, i.e., if learner u answers e_i correctly, $y_i = 1$ otherwise $y_i = 0$. c_e denotes textual information of each exercise $e \in E$, e.g., textual content and corresponding knowledge concepts. We provide c_e in a $\langle key, value \rangle$ form, as the example in Figure 1.

Based on the above conditions, the simulator’s overarching goal is to accurately generate the human learners’ future response data on unseen exercises by faithfully distilling their learning patterns and cognitive preferences from historical response data. Please note that existing personalized learning algorithms usually assume that learners only submit each exercise once, so repeated submission is not considered in our simulation.

3.2 LLM-powered Agent

The generative agent in Agent4Edu uses LLM as its foundational architecture, enhancing its functionality tailored for the personalized learning scenario through three specialized modules: learner profile, memory, and action modules. To mimic actual personalized practice responses akin to humans, we construct an individual agent $agent_u$ for each learner u . Each agent integrates a learner profile module aimed at reflecting personalized practice patterns and cognitive factors. Additionally, each agent is equipped with a memory module designed to store past practice records and summarize high-level ideas. To simulate learner practice behavior more cohesively, the agent is also equipped with an action module.

Learner Profile Module The learner profile module represents some overall learning features of human learners, which are typically stable and derived from long-term learning experiences. We configure each agent $agent_u$ ’s profile based on its corresponding learner u ’s response data². Each agent’s initial configuration is divided into two categories: explicit practice styles and implicit cognitive factors.

Practice styles are statistical features explicitly derived from the available practice record l_u of each learner u , such as learning activity (Baker 2001; Gao et al. 2021), practice diversity (Bi et al. 2020), success rate, and preference. Activity indicates learners’ enthusiasm for learning and provides clues for simulating their practice behaviors. For example, learners with higher enthusiasm for learning usually perform better. Mathematically, the activity level of learner u is defined as $P_{act}^u = \frac{|l_u|}{|E|}$. Practice diversity reflects the knowledge coverage practiced by learners, represented as $P_{div}^u = \frac{|K_u|}{|K|}$, where $|K_u|$ is the number of knowledge concepts practiced by learner u . Higher diversity indicates greater curiosity in learners. Success rate correlates with the probability of learners answering questions correctly, making it another essential characteristic. The success rate for learner u is mathematically represented as $P_{suc}^u = \frac{\sum_{y_{u,e_i} \in l_u} y_{u,e_i}}{|l_u|}$. Preference refers to the knowledge concepts that learners practice most frequently.

Cognitive factors are implicit features studied in psychology (Baker 2001; Chen et al. 2024), which significantly impact learner u ’s practice performance. We select problem-solving ability and knowledge proficiency (Cheng et al. 2024) for this study. Problem-solving ability is assumed to be stable during the learning process, while knowledge proficiency typically evolves with learning progress (Huang et al. 2020). Therefore, in the profile module, we only configure the ability factor P_{ab}^u , with knowledge mastery being considered in the subsequent memory module. To obtain implicit ability, we assign a psychological IRT model (Baker 2001) trained on the observed learner response records, as the tool for the agent, allowing it to infer each learner u ’s ability factor from the response data l_u . The training and use of the IRT tool are detailed in Appendix B.

²Note that if zero-shot simulations are performed and user data is unavailable, the profile needs to be randomly generated.

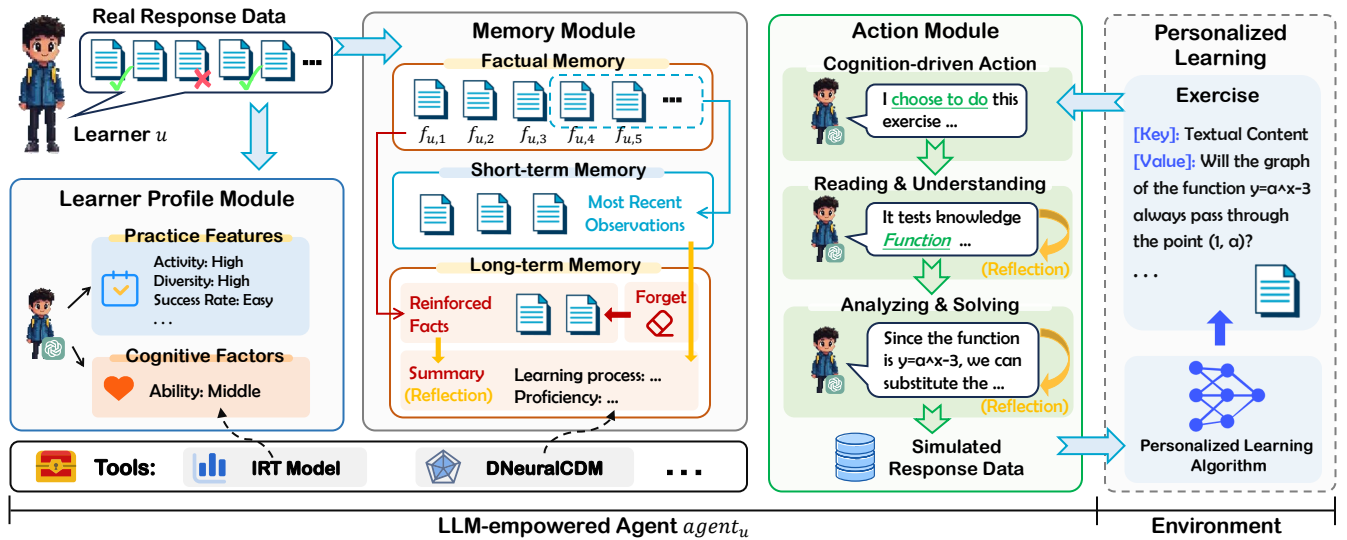


Figure 1: The overall framework of Agent4Edu.

Notably, we segment the values of each of the above features into several tiers in order to better prompt the generative agent inspired by (Wang et al. 2023d). For a detailed exposition, refer to Appendix A.1. Additionally, to ensure broad applicability and protect privacy, certain personal identifiers (such as name, gender, age, and occupation) are intentionally anonymized in this work (Zhang et al. 2023a; Li et al. 2023). While these attributes may help shape other types of agents, they are not primary factors affecting practice performance in education. Our approach, based on both behavioral practice styles and psychological cognitive settings, can support a comprehensive representation of real learners.

Memory Module The Memory module allows the LLM-based agent $agent_u$ to observe and summarize its corresponding learner u 's past practice experiences step by step. This module provides insightful clues to the agent for response simulation on unseen exercises. We follow the human learning mechanism (Atkinson 1968a; Cowan 2008; Huang et al. 2020; Wang et al. 2023d) to design three types of memories for each agent: *factual memory*, *short-term memory*, and *long-term memory*. Each memory is initially set to empty.

- **Factual Memory:** In our simulation, factual memory is defined as the true learner's past response records (i.e., *observations*). When the agent obtains a new response record of learner u at step i , i.e., $l_{u,i} = (e_i, c_{e_i}, y_{u,e_i})$, the response record is transmitted to the factual memory for processing.

Inspired by human learning mechanisms, if an agent repeatedly practices similar questions or knowledge, their memory is strengthened (Huang et al. 2020). Therefore, we introduce an additional counter $f_{u,i}$ (initially set to 1) for each record $l_{u,i}$ in factual memory to track the number of times it has been reinforced, a simple yet effective method that has been successfully used in user preference simulation (Wang et al. 2023d). Formally, for each

agent $agent_u$, assume it has observed n factual memories is $M_u = \{l_{u,1}, l_{u,2}, \dots, l_{u,n}\}$, then it is allowed to receive a new response record $l_{u,n+1}$. We first calculate the similarity between $l_{u,n+1}$ and each existing factual memory $l_{u,i}$ in the current memory M_u . The similarity between records can be defined as a metric that can be evaluated by LLMs, cosine similarity between text vectors, and other similar measures. In this case, we use the similarity relationships between the knowledge concepts involved in the records for the calculation. Specifically, we employ the statistical tool³ released by RCD (Gao et al. 2021) to determine whether two knowledge concepts are similar. If there is a similarity between the knowledge concepts involved in two records, the two records are considered similar. For similar records, we increment the counter for $l_{u,i}$ by 1 (i.e., $f_{u,i} \leftarrow f_{u,i} + 1$), indicating that it has been reinforced by $l_{u,n+1}$, and then add $l_{u,n+1}$ to factual memory; otherwise, $l_{u,n+1}$ is directly recorded without any reinforcement. After processing and saving a new response record, factual memory triggers updating *short-term* and *long-term* memories.

We emphasize that the agent can only save response records into factual memory but cannot directly retrieve it, thereby allowing the retention of all exercise textual information and responses without being constrained by the LLM's context length limitations.

- **Short-term Memory:** Human short-term memory refers to the recent and temporary information that can be retained and recalled over a relatively brief period (Atkinson 1968b). Therefore, in our simulation, short-term memory is employed to retain the details of the agent's most recent observed s records. Assuming the current factual memory of agent $agent_u$ is $M_u = \{l_{u,1}, l_{u,2}, \dots, l_{u,n}\}$, then the short-term memory storage is defined as $M_{u,short} = \{l_{u,n-s+1}, \dots, l_{u,n}\}$.

- **Long-term Memory:** Long-term memory is formed through

³<https://github.com/bigdata-ustc/RCD>

the reinforcement of memories from repeated practice and self-reflection inspired by to human long-term memories (Matelsky et al. 2023). It possesses a wide receptive field, allowing it to retain information observed long ago and generate high-level insights. We design the long-term memory using three types of information: (1) **Reinforced Facts**: During each update of long-term memory, the agent $agent_u$ first goes through the current factual memory M_u . When the count $f_{u,i}$ of a record $l_{u,i}$ exceeds a preset threshold F , indicating that the memory has been reinforced F times, it is converted into long-term memory. (2) **Learning Process Summary**: We utilize the LLM embedded in the agent to summarize the agent’s learning status from both short-term and long-term memories by *Memory Reflection*. Each step of the summary replaces the previous summary. The summary consists of linguistic descriptions of the practice process and new insights from the agent itself. It overlooks practice details to filter out noise, irrelevant content, or potentially misleading information. Furthermore, compressing memory conserves significant space and enhances operational efficiency. (3) **Knowledge Proficiency**: We allow the agent to use an optimized DNeuralCDM (Wang et al. 2023a) based on the observed learner response data as a tool to obtain the learner’s dynamic proficiency (segmented into several tiers) evolution of specific knowledge concepts after each step of practice. The knowledge proficiency is a kind of dynamic cognitive factor significantly reflecting human responses in education (Piech et al. 2015; Wang et al. 2024a). The training and use of DNeuralCDM are given in Appendix B.

Additionally, each factual record in long-term memory may be **forgotten** following the human forgetting curve theory (Averell and Heathcote 2011; Huang et al. 2020) that human memory decay starts rapidly and then gradually slows over time. We define a forgetting function associated timestamp i and current observed step n , i.e., $g(l_{u,i}) = \frac{1}{1+e^{-(n-i)}}$, to simulate human learners’ forgetting. For each factual record in the long-term memory M_u , it is forgotten if $g(l_{u,i})$ exceeds a predetermined threshold λ and its reinforcement frequency $f_{u,i}$ in factual memory is then reset as 1.

Overall, the factual response records are specific, while learning memory summaries are more general. By combining them, the agent can accurately perceive the learner’s practice process. Please note that traditional simulators (Piech et al. 2015; Zhao et al. 2023) can be regarded as owning the short-term memory but no long-term memory.

To help agents interact with the personalized learning environment, we introduce three memory operations:

- *Memory Retrieval*: This operation assists the agent in extracting related information from memory. We allow the agent to retrieve the short-term and long-term memories finding reinforced facts and conducting summary.
- *Memory Writing*: The raw observations are firstly input into the factual memory as facts. Then the recent facts are stored in short-term memory and the reinforced facts are written into long-term memory.
- *Memory Reflection*: This operation occurs exclusively within long-term memory containing two aspects of reflections: (1) **Summary Reflection** is performed to summarize

high-level ideas based on short-term and long-term memories, and (2) **Corrective Reflection** is performed when the agent’s action is inconsistent with the real learner, which will be introduced in Action Module.

Action Module To equip the agent with learner profiles and memory modules, enabling it to exhibit human-like problem-solving behaviors and responses based on current observations, we design a specialized action module for each agent within Agent4Edu tailored for personalized learning. This module encompasses three main categories of actions:

- *Cognition-driven Actions*: In our simulation, personalized learning algorithms recommend one exercise to the agent at each step. The agent read the exercise’s content and decides whether or not to practice it, based on current cognitive factors. If the exercise is too challenging relative to the agent’s assessed ability and knowledge proficiency, the agent can opt to reject the recommended exercise.

- *Reading and Understanding Exercises*. Simulating the process of reading and understanding exercises, similar to how humans approach them, provides valuable and interpretable insights into the agents. During each practice session, the agent is first required to identify and describe a knowledge concept assessed by the current exercise. If the agent correctly matches the exercise’s knowledge concept, it demonstrates an understanding of the exercise context akin to human learners. If the agent fails to do so, a corrective reflection is triggered to guide the agent towards the correct knowledge concept. This method reduces the risk of inaccuracies and ensures the agent’s credibility in simulating learner response (Zhang et al. 2023a).

- *Analyzing and Solving Exercises*. Analyzing and solving exercises are crucial aspects of the learning process. Unlike previous simulation methods that directly predict the learner’s response in terms of answer correctness, our simulation requires the agent to emulate the learner’s answering process, which enhances both interpretability and credibility. To simulate this complex answer process more effectively, we improve agent’s reasoning ability through a chain-of-thought approach (Wei et al. 2022). Initially, the agent combines its profile and memories to formulate an initial solution idea for the exercise. Then, it writes the final answer to the exercise based on the solution idea. Afterwards, the agent predicts whether its answer is correct (i.e., performance prediction). If the predicted response does not match the real learner’s response, a corrective reflection is triggered. Note that, if standard answers of exercises are available, a scoring program can be designed to directly assess the correctness of the agent’s answer.

3.3 Personalized Learning Scenarios

Agent4Edu simulates agent and learning environment interaction (see Appendix D for a case study). The learning environment is designed as a standalone module that incorporates a series of personalized algorithms. These algorithms can recommend exercises to agents based on their past practice data. For instance, our experiments utilize computerized adaptive testing (CAT) strategies (Bi et al. 2020) for personalized learning. The module features an open interface,

allowing researchers and practitioners to integrate external personalized learning algorithms seamlessly. This adaptability ensures that Agent4Edu serves as a versatile platform for comprehensive evaluations and the future collection of valuable learner response data.

4 Experiment

Dataset Our dataset, called *EduData*, is provided by iFLYTEK Co., Ltd. It comprises 18,045 time-ordered response records from 500 Chinese high school students in the subjects of mathematics and physics. Each record includes the exercise ID, correctness, and timestamp. There are 1,032 exercises and 458 knowledge concepts in total, with each exercise testing one knowledge concept. Additionally, to facilitate reasoning and reflection for LLM-based agents, the platform provider has furnished us with the textual content of the exercises. In the experiment, we translate all Chinese text of exercises into English.

Experimental Setup We use GPT-3.5-turbo and GPT-4 through OpenAI’s API service ⁴ to construct the agent for experimentation. When operating under the GPT-3.5-turbo configuration, all response data is utilized for experiments. Due to cost considerations, we simulate the task records of only 100 learners under the GPT-4 setting. The temperature parameter of GPT is 0 to avoid randomness. Empirically, we set the short-term memory size to 5, the threshold F for memory enhancement is 5, and the threshold λ for forgetting in long-term memory is 0.99. Note that, in our experiments, unless explicitly specified, the LLM used is GPT-3.5-turbo.

4.1 LLM-based Agent Simulation Evaluation

Motivation: The LLM-based agent is the core component of Agent4Edu. Exploring whether the agent can truly simulate human learners’ practice response is crucial for enhancing intelligent educational systems. We evaluate the effectiveness of the generative agent, including response simulation, exercises’ knowledge understanding, zero-shot simulations, and ablation experiments.

Learner Simulation Evaluation The agent aims to generate simulated learner response data that closely approximates real responses. To validate the effectiveness of the simulation, we compare it with two traditional supervised simulation methods, including DAISIM (Zhao et al. 2023) and KES (Liu et al. 2019). Additionally, to enrich our baseline for a compelling comparison, we include several Knowledge Tracing (KT) models, such as DKVMN (Zhang et al. 2017), EERNN (with Markov) (Su et al. 2018) and SAKT (Pandey and Karypis 2019), which are similar to the learner simulator in terms of response prediction.

In the experimental setup, each learner’s records are divided into a 90% training set and a 10% test set. Each baseline model which is data-driven is trained on the training data, with the last 20% records of each learner’s training data used for model validation. The agent has access to all training data to generate profiles and update its memory through

⁴The detailed GPT versions: GPT-3.5-turbo-1106 (up to Sep. 2021) and GPT-4-turbo (up to Dec. 2023).

Model	ACC \uparrow	F1-score \uparrow	ROUGE-3 \uparrow
KES	50.11	58.32	25.77
DKVMN	64.39	76.70	37.24
EERNN	<u>65.72</u>	76.06	43.55
SAKT	65.52	<u>78.33</u>	31.09
DAISIM	65.63	78.25	31.72
Agent4Edu (GPT-3.5-turbo)	66.70	79.84	<u>37.97</u>
Agent4Edu (GPT-3.5-turbo) ₁₀₀	65.40	78.72	35.14
Agent4Edu (GPT-4) ₁₀₀	66.51	79.53	34.86

Table 1: Prediction scores (%) on evaluating simulation performance. The best results are bold, the second-best results are marked by an underline, and \uparrow means the higher score the better performance, the same as below. Agent4Edu₁₀₀ indicates a basic exploratory on simulating 100 learners.

reflection. During the testing phase, each trained baseline is tasked with predicting learners’ binary responses (correct or incorrect) to unseen exercises in the test data. For our generative agent, exercises from the test data are sequentially sent to it, and it performs the designed three actions to solve them. If the agent rejects an exercise due to its difficulty, we label its response as an “incorrect answer”. The evaluation metrics are selected from two perspectives. Firstly, we use accuracy (ACC) and F1-score to measure prediction accuracy. Secondly, we assess the similarity between the simulated and real data distributions using ROUGE-3, inspired by (Zhao et al. 2023). We repeatedly run each baseline model five times in the same setups and the Table 1 reports the average scores.

The experimental results indicate that Agent4Edu (GPT-3.5-turbo) demonstrates strong competitiveness compared to the supervised baselines, particularly in terms of ACC and F1-score. This suggests that the LLM-based agent has the potential to generate learner response data that closely resembles real-world datasets. Furthermore, among the baselines, EERNN performs exceptionally well by effectively modeling the exercise content as supplementary clues. Finally, an exploratory simulation conducted using Agent4Edu (GPT-3.5-turbo and GPT-4) on a subset of data with 100 learners shows that they enable the simulated distribution to closely approach the real distribution. Among them, GPT-4 performs better in terms of ACC and F1-score.

Additionally, we evaluate whether the simulated distribution of the agent’s practice success rate aligns with the actual distribution of learner data. We use the real response success rate as the ground truth and then replace the corresponding responses in the real sequence with the predicted responses from the test data to calculate the agent’s simulated success rate, as shown in Figure 2 (a). The comparison between the ground truth values and the agent’s results indicates that the simulated data effectively captures the learners’ practice patterns related to success rate.

Understanding Exercise-related Knowledge To evaluate whether the agent understands a specific exercise, the agent is tasked with generating the knowledge concept tested by the exercise. Specifically, we create a candidate list containing one actual knowledge concept related to the exercise

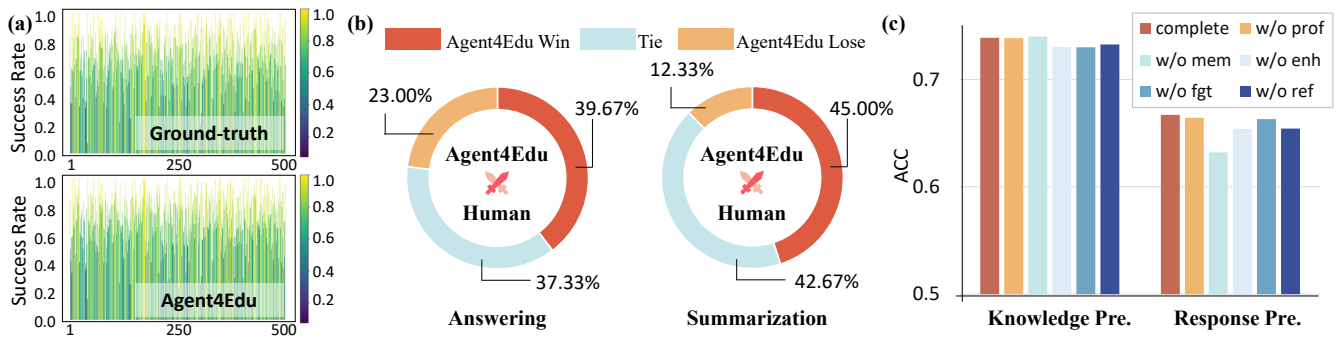


Figure 2: (a) Comparison between the success rate distributions of ground-truth and agent-simulated response data. (b) Using LLMs as judges to identify whether the records of the agent’s simulations originate from real humans. (c) The ablation studies.

Model	ACC \uparrow
Agent4Edu (GPT-3.5-turbo)	73.88
Agent4Edu (GPT-3.5-turbo) ₁₀₀	74.57
Agent4Edu (GPT-4) ₁₀₀	82.43

Table 2: The ACC of knowledge prediction.

and two random knowledge concepts unrelated to the exercise. The agent must then select the relevant knowledge concept from this list based on its understanding (detailed prompts are provided in Appendix C). We use ACC as the metric to evaluate the agent’s knowledge predictions for all exercises in the test set, treating it as a binary classification task. This section uses the same agents from the section “Learner Simulation Evaluation”. The experimental results presented in Table 2, indicate that all the agents can correctly identify the knowledge being tested in most practice exercises. This demonstrates the strong human-like ability and rich knowledge of LLMs to comprehend exercises. Furthermore, under the same conditions with 100 learners, the agent with GPT-4 is more accurate than the one with GPT-3.5-turbo, indicating that GPT-4 has a stronger semantic understanding ability compared to GPT-3.5-turbo.

Zero-shot Simulation Zero-shot simulation presents a significant challenge in real-world applications, particularly when learners are in a cold-start situation where their response data is unavailable. This limitation restricts the applicability of previous simulation models. To validate the zero-shot simulation capability of the agent, we initialize 10 agents with randomly generated profiles and have them sequentially answer 10 randomly selected exercises with IDs {120, 250, 113, 1330, 568, 1881, 771, 593, 1, 595}. In this zero-shot scenario, we disable the corrective reflection mechanism and tools, due to the absence of learner response data. The summary reflection remains usable for the agent. Three GPT-3.5-turbo models, with a temperature parameter of 0.5, act as annotators, tasked with evaluating whether each simulated record (including exercise answers and practice summaries) is written by a real human. Records deemed to be human-written are labeled as “Agent4Edu

Win”, non-human records are labeled as “Lose”, and ambiguous records are labeled as “Tie”.

The results depicted in Figure 2 (b) indicate that the agent’s performance in summarization is closely aligned with the real human responses, making differentiation between the two challenging. However, the agent exhibits certain limitations in answering exercise tasks compared to summarization tasks, primarily due to the complexity of reasoning required to solve exercises.

Ablation Study We conduct ablation studies to evaluate the impact of key components within the GPT-3.5-turbo-powered agent. The results illustrated in Figure 2 (c) show the accuracy of the agent’s exercise-related knowledge prediction and response prediction under various conditions: without the profile module (w/o prof), without the memory module (w/o mem), without the memory enhancement (w/o enh), without the memory forgetting (w/o fgt), and without reflection (w/o ref). These findings confirm the effectiveness of each component in improving the agent’s predictive performance on learners’ response data. However, the ablation experiments indicate that the impact on knowledge prediction is not significant. This can be attributed to the fact that the original GPT-3.5-turbo model already possesses a substantial amount of knowledge, which is sufficient to support exercise comprehension.

4.2 Personalized Learning

Motivation The primary objective of Agent4Edu is to comprehensively and accurately evaluate personalized learning algorithms and use the generated data to enhance their effectiveness. We aim to validate this objective from two perspectives: (1) through an agent-based multifaceted evaluation of personalized learning services, and (2) by assessing the potential improvements in personalized learning algorithms based on the simulated data.

Multifaceted Evaluation Human learners have multifaceted evaluations of different personalized learning services, such as whether the recommended task difficulty is too challenging. Assuming that a generative agent can accurately simulate the behavior of real learners, its evaluation of personalized algorithms tends to align with human evaluations. We utilize the Computerized Adaptive Testing

Model	satisfaction	AoD	gain
FSI	<u>39</u>	70	48
KLI	<u>39</u>	66	43
MAAT	42	<u>68</u>	<u>45</u>

Table 3: Multifaceted evaluations of CAT strategies.

(CAT) which aims to estimate learners’ ability or knowledge proficiency with minor exercises, as the experimental environment, including FSI (Lord 2012), KLI (Chang and Ying 1996), and MAAT (Bi et al. 2020). We use 100 randomly initialized agents to generate virtual data for pretraining the cognitive diagnosis model (i.e., the IRT model (Baker 2001)) in the CAT algorithm for learner evaluation. Based on this, each adaptive algorithm iterates through 10 rounds to recommend 100 randomly initialized agents (zero-shot simulation), with one exercise recommendation per round. Upon the conclusion of personalized testing, each agent is required to evaluate each CAT algorithm. To achieve this, we design three evaluation metrics, including satisfaction, appropriateness of difficulty (AoD), and whether there was any gain. Table 3 presents a comprehensive evaluation of various strategies, where the element in the i -th row and j -th column represents the number of agents that consider the corresponding CAT algorithm i to meet the metric j . Clearly, the agent demonstrates higher satisfaction in recommending MAAT. This observation aligns with the common understanding in the research community that MAAT considers both the difficulty of exercises and the diversity of knowledge (Bi et al. 2020), making the overall service more reasonable. Additionally, FSI focuses on recommending exercises that are moderately difficult and likely to provide gain. These findings highlight the LLM-powered agent’s fine-grained evaluation level for learning algorithms.

Personalized Learning Algorithm Improvement We investigate whether the simulated data generated by Agent4Edu can enhance personalized learning algorithms. We select CAT as our personalized learning assessment task due to their representativeness in intelligent education. If the generated data can improve the performance of CAT models, it will indicate the effectiveness of our proposed Agent4Edu.

To set up, we select 60% of the learners’ data from EduData to train the cognitive diagnosis models (i.e., the IRT model) for learner evaluation in CAT algorithms (i.e., FSI (Lord 2012), KLI (Chang and Ying 1996), and MAAT (Bi et al. 2020)). The remaining 40% of learners’ data is used to test the CAT models. Furthermore, for each learner in the test data, we simulate their responses to 20 randomly selected unseen exercises based on their profiles. Using this strategy, we generate simulated learner data, which are then merged with the training data from the original EduData to form the augmented dataset, EduData+. We train the IRT model in each CAT model using both the original EduData and EduData+, and then evaluate each CAT strategy by recommending 5 and 10 test exercises for each learner.

Table 4 lists the IRT prediction performance after retraining on the testing records via CAT, where F1-score repre-

Testing length is 5			
Model	F1-score	F1-score+	Imp.
FSI	80.11	82.39	+2.28
KLI	79.45	81.84	+2.39
MAAT	81.77	81.97	+0.20
Testing length is 10			
Model	F1-score	F1-score+	Imp.
FSI	81.10	82.51	+1.41
KLI	80.63	82.82	+2.19
MAAT	81.71	81.88	+0.17

Table 4: The improvement of CAT services.

sents scores on EduData, and F1-score+ represents scores on EduData+. The results demonstrate that CAT strategies can be effectively enhanced with the assistance of Agent4Edu. This suggests that Agent4Edu is capable of generating high-quality learner response data, even with randomly initialized agents (in zero-shot scenarios), thereby enriching the provided dataset.

5 Conclusion

In this paper, we introduce Agent4Edu, an innovative personalized learning simulator that leverages LLM-powered generative agents to simulate learners’ response data, as well as detailed problem-solving behaviors. Our generative agents are equipped with learner Profile, Memory and Action modules specifically tailored for personalized learning scenarios. These agents exhibit human-like choosing, understanding, analyzing and answering exercises, which accurately predict their future responses. Additionally, the generative agent can interact with personalized learning environments to evaluate and enhance intelligent services. Through comprehensive and meticulous evaluation, we explore the strengths and weaknesses of Agent4Edu, emphasizing the consistency and discrepancies in practice behaviors observed between agents and learners. In the future, we plan to research multi-learner agent cooperation and multi-modal practice solutions using generative agents. We hope that our research will provide new insights into the field of intelligent education.

Acknowledgements

This research was partially supported by grants from the National Natural Science Foundation of China (No.62337001, 62477044), the Key Technologies R & D Program of Anhui Province (No. 202423k09020039), and the Fundamental Research Funds for the Central Universities.

References

- Atkinson, R. C. 1968a. Human memory: A proposed system and its control processes. *The Psychology of Learning and Motivation*, 2.
- Atkinson, R. C. 1968b. A proposed system and its control processes. *The Psychology of Learning and Motivation*, 2.

- Averell, L.; and Heathcote, A. 2011. The form of the forgetting curve and the fate of memories. *Journal of mathematical psychology*, 55(1): 25–35.
- Baidoo-Anu, D.; and Ansah, L. O. 2023. Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI*, 7(1): 52–62.
- Baker, F. B. 2001. *The basics of item response theory*. ERIC.
- Bi, H.; Ma, H.; Huang, Z.; Yin, Y.; Liu, Q.; Chen, E.; Su, Y.; and Wang, S. 2020. Quality meets diversity: A model-agnostic framework for computerized adaptive testing. In *2020 IEEE International Conference on Data Mining (ICDM)*, 42–51. IEEE.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chang, H.-H.; and Ying, Z. 1996. A global information approach to computerized adaptive testing. *Applied Psychological Measurement*, 20(3): 213–229.
- Chen, X.; Wu, L.; Liu, F.; Chen, L.; Zhang, K.; Hong, R.; and Wang, M. 2024. Disentangling Cognitive Diagnosis with Limited Exercise Labels. *Advances in Neural Information Processing Systems*, 36.
- Cheng, K.; Peng, L.; Wang, P.; Ye, J.; Sun, L.; and Du, B. 2024. DyGKT: Dynamic Graph Learning for Knowledge Tracing. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 409–420.
- Cowan, N. 2008. What are the differences between long-term, short-term, and working memory? *Progress in brain research*, 169: 323–338.
- Dan, Y.; Lei, Z.; Gu, Y.; Li, Y.; Yin, J.; Lin, J.; Ye, L.; Tie, Z.; Zhou, Y.; Wang, Y.; et al. 2023. Educhat: A large-scale language model-based chatbot system for intelligent education. *arXiv preprint arXiv:2308.02773*.
- Gao, C.; Lan, X.; Lu, Z.; Mao, J.; Piao, J.; Wang, H.; Jin, D.; and Li, Y. 2023. S³: Social-network Simulation System with Large Language Model-Empowered Agents. *arXiv preprint arXiv:2307.14984*.
- Gao, W.; Liu, Q.; Huang, Z.; Yin, Y.; Bi, H.; Wang, M.-C.; Ma, J.; Wang, S.; and Su, Y. 2021. RCD: Relation map driven cognitive diagnosis for intelligent education systems. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, 501–510.
- Huang, X.; Lian, J.; Lei, Y.; Yao, J.; Lian, D.; and Xie, X. 2023. Recommender ai agent: Integrating large language models for interactive recommendations. *arXiv preprint arXiv:2308.16505*.
- Huang, Z.; Liu, Q.; Chen, Y.; Wu, L.; Xiao, K.; Chen, E.; Ma, H.; and Hu, G. 2020. Learning or forgetting? a dynamic approach for tracking the knowledge proficiency of students. *ACM Transactions on Information Systems (TOIS)*, 38(2): 1–33.
- Huang, Z.; Liu, Q.; Zhai, C.; Yin, Y.; Chen, E.; Gao, W.; and Hu, G. 2019. Exploring multi-objective exercise recommendations in online education systems. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 1261–1270.
- Jin, D.; Mehri, S.; Hazarika, D.; Padmakumar, A.; Lee, S.; Liu, Y.; and Namazifar, M. 2023. Data-efficient alignment of large language models with human feedback through natural language. *arXiv preprint arXiv:2311.14543*.
- Kieser, F.; Wulff, P.; Kuhn, J.; and Küchemann, S. 2023. Educational data augmentation in physics education research using ChatGPT. *Physical Review Physics Education Research*, 19(2): 020150.
- Li, H.; Xu, T.; Zhang, C.; Chen, E.; Liang, J.; Fan, X.; Li, H.; Tang, J.; and Wen, Q. 2024. Bringing generative AI to adaptive learning in education. *arXiv preprint arXiv:2402.14601*.
- Li, Y.; Chen, X.; Zhao, H.; Gong, J.; Zhou, G.; Rossano, F.; and Zhu, Y. 2023. Understanding Embodied Reference with Touch-Line Transformer. In *ICLR*.
- Liu, J.; Huang, Z.; Xiao, T.; Sha, J.; Wu, J.; Liu, Q.; Wang, S.; and Chen, E. 2024. SocraticLM: Exploring Socratic Personalized Teaching with Large Language Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Liu, Q.; Tong, S.; Liu, C.; Zhao, H.; Chen, E.; Ma, H.; and Wang, S. 2019. Exploiting cognitive structure for adaptive learning. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 627–635.
- Liu, R.; Yang, R.; Jia, C.; Zhang, G.; Zhou, D.; Dai, A. M.; Yang, D.; and Vosoughi, S. 2023. Training socially aligned language models in simulated human society. *arXiv preprint arXiv:2305.16960*.
- Long, X.; Zeng, J.; Meng, F.; Ma, Z.; Zhang, K.; Zhou, B.; and Zhou, J. 2024. Generative multi-modal knowledge retrieval with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18733–18741.
- Lord, F. M. 2012. *Applications of item response theory to practical testing problems*. Routledge.
- Matelsky, J. K.; Parodi, F.; Liu, T.; Lange, R. D.; and Kording, K. P. 2023. A large language model-assisted education tool to provide feedback on open-ended responses. *arXiv preprint arXiv:2308.02439*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Pandey, S.; and Karypis, G. 2019. A self-attentive model for knowledge tracing. In *12th International Conference on Educational Data Mining, EDM 2019*, 384–389. International Educational Data Mining Society.
- Park, J. S.; O’Brien, J.; Cai, C. J.; Morris, M. R.; Liang, P.; and Bernstein, M. S. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the*

36th Annual ACM Symposium on User Interface Software and Technology, 1–22.

Piech, C.; Bassen, J.; Huang, J.; Ganguli, S.; Sahami, M.; Guibas, L. J.; and Sohl-Dickstein, J. 2015. Deep knowledge tracing. *Advances in neural information processing systems*, 28.

Qadir, J. 2023. Engineering education in the era of ChatGPT: Promise and pitfalls of generative AI for education. In *2023 IEEE Global Engineering Education Conference (EDUCON)*, 1–9. IEEE.

Qian, C.; Cong, X.; Yang, C.; Chen, W.; Su, Y.; Xu, J.; Liu, Z.; and Sun, M. 2023. Communicative agents for software development. *arXiv preprint arXiv:2307.07924*.

Rahman, M. M.; and Watanobe, Y. 2023. ChatGPT for education and research: Opportunities, threats, and strategies. *Applied Sciences*, 13(9): 5783.

Reddy, S.; Levine, S.; and Dragan, A. 2017. Accelerating human learning with deep reinforcement learning. In *NIPS workshop: teaching machines, robots, and humans*.

Su, Y.; Liu, Q.; Liu, Q.; Huang, Z.; Yin, Y.; Chen, E.; Ding, C.; Wei, S.; and Hu, G. 2018. Exercise-enhanced sequential modeling for student performance prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.

Wang, F.; Gao, W.; Liu, Q.; Li, J.; Zhao, G.; Zhang, Z.; Huang, Z.; Zhu, M.; Wang, S.; Tong, W.; et al. 2024a. A Survey of Models for Cognitive Diagnosis: New Developments and Future Directions. *arXiv preprint arXiv:2407.05458*.

Wang, F.; Huang, Z.; Liu, Q.; Chen, E.; Yin, Y.; Ma, J.; and Wang, S. 2023a. Dynamic cognitive diagnosis: An educational priors-enhanced deep knowledge tracing perspective. *IEEE Transactions on Learning Technologies*, 16(3): 306–323.

Wang, G.; Xie, Y.; Jiang, Y.; Mandlekar, A.; Xiao, C.; Zhu, Y.; Fan, L.; and Anandkumar, A. 2023b. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.

Wang, L.; Ma, C.; Feng, X.; Zhang, Z.; Yang, H.; Zhang, J.; Chen, Z.; Tang, J.; Chen, X.; Lin, Y.; et al. 2024b. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6): 1–26.

Wang, L.; Zhang, J.; Chen, X.; Lin, Y.; Song, R.; Zhao, W. X.; and Wen, J.-R. 2023c. Recagent: A novel simulation paradigm for recommender systems. *arXiv preprint arXiv:2306.02552*.

Wang, L.; Zhang, J.; Yang, H.; Chen, Z.; Tang, J.; Zhang, Z.; Chen, X.; Lin, Y.; Song, R.; Zhao, W. X.; et al. 2023d. User behavior simulation with large language model based agents. *arXiv preprint arXiv:2306.02552*.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.

Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Zhang, S.; Zhu, E.; Li, B.; Jiang, L.; Zhang, X.; and Wang, C. 2023. Autogen:

Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*.

Xu, S.; Zhang, X.; and Qin, L. 2024. EduAgent: Generative Student Agents in Learning. *arXiv preprint arXiv:2404.07963*.

Yao, F.; Liu, Q.; Yue, L.; Gao, W.; Li, J.; Li, X.; and He, Y. 2024. Adard: An adaptive response denoising framework for robust learner modeling. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3886–3895.

Yue, L.; Liu, Q.; Du, Y.; Gao, W.; Liu, Y.; and Yao, F. 2023. Fedjudge: Federated legal large language model. *arXiv preprint arXiv:2309.08173*.

Zhang, A.; Sheng, L.; Chen, Y.; Li, H.; Deng, Y.; Wang, X.; and Chua, T.-S. 2023a. On generative agents in recommendation. *arXiv preprint arXiv:2310.10108*.

Zhang, J.; Hou, Y.; Xie, R.; Sun, W.; McAuley, J.; Zhao, W. X.; Lin, L.; and Wen, J.-R. 2023b. Agentcf: Collaborative learning with autonomous language agents for recommender systems. *arXiv preprint arXiv:2310.09233*.

Zhang, J.; Shi, X.; King, I.; and Yeung, D.-Y. 2017. Dynamic key-value memory networks for knowledge tracing. In *Proceedings of the 26th international conference on World Wide Web*, 765–774.

Zhao, G.; Huang, Z.; Zhuang, Y.; Liu, J.; Liu, Q.; Liu, Z.; Wu, J.; and Chen, E. 2023. Simulating Student Interactions with Two-stage Imitation Learning for Intelligent Educational Systems. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 3423–3432.

Zhuang, Y.; Liu, Q.; Zhao, G.; Huang, Z.; Huang, W.; Pardos, Z.; Chen, E.; Wu, J.; and Li, X. 2024. A Bounded Ability Estimation for Computerized Adaptive Testing. *Advances in Neural Information Processing Systems*, 36.