Strategy-based Classroom Simulation: Enhancing Cognitive Elicitation in LLM Classroom Agents

Anonymous ACL submission

Abstract

Large Language Models (LLMs) in education have demonstrated satisfactory performance in classroom simulations and can generate rich conversational content. However, the cognitive guidance functions behind such dialogues in real classrooms have not been precisely modeled. To address this gap, we propose **Strategy-based Class Simulation (SCS)**, a simulation framework that enhances the agent's cognitive elicitation ability. Specifically, we introduce two methods to better align simulated data with real-world data: **Strategy Recommendation**, based on strategic distribution information derived from authentic classroom interactions, and **ε-greedy Strategy Selection (ε-GSS)**, which enables richer agent choices. Building on these methods, we construct a classroom simulation task set, SCST-100, with real classroom data. Results from simulation experiments show that SCS agents achieve greater fidelity compared to direct simulation.

1 Introduction

Recent years have seen the development of Large Language Model (LLM) agents for educational scenario simulation (Pedersen and Duin, 2022; Dai et al., 2024; Jinxin et al., 2023). These agents use environmental cues and simulation frameworks to generate highly human-like action trajectories and dialogues. As a result, they achieve strong coherence and accuracy (Zhang et al., 2025; Jin et al., 2025).

Most existing work focuses on knowledge transmission in classrooms. Agents act as either teachers or assistants to provide instruction (Tu et al., 2023; Sonkar et al., 2023). They can also be students who acquire knowledge (Lim et al., 2025; Markel et al., 2023). However, classroom dialogue is a critical tool for cognitive development, not just for delivering information. Beyond knowledge, dialogue involves questioning, responding,

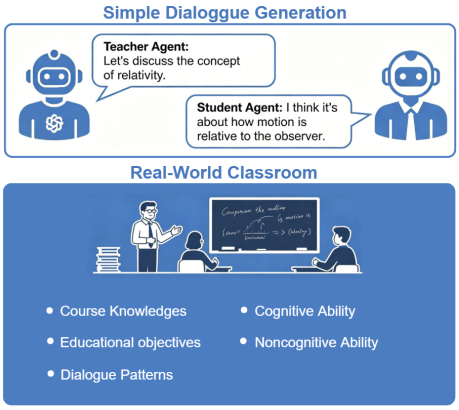

Figure 1: Compared with simple dialogue generation, real classroom dialogue encompasses multi-dimensional information.

and guiding. These processes stimulate cognitive ability (Howe et al., 2019). The difference in information content between simple simulations and real scenarios is significant, as shown in Figure 1. Relying solely on textual dialogue limits insights into classroom teaching paradigms.

To incorporate cognitive information into dialogue, we introduce a dialogue strategy system that categorizes dialogues based on their cognitive development roles into strategy groups. This system is central to our **Strategy-based Class Simulation (SCS)** framework, which models teacher-student dialogue patterns. It enables agents to select strategies autonomously during interaction. However, relying solely on basic strategy selection does not accurately reflect real-world dialogue. To address this issue, we adopt two techniques:

- **Strategy Recommendation** We construct a recommendation system that leverages real-world data statistics to suggest dialogue strategies based on conversation history, thereby enhancing the rationality of strategy selection.
- **ε-greedy Strategy Selection (ε-GSS)** We employ ε-greedy in strategy selection, making

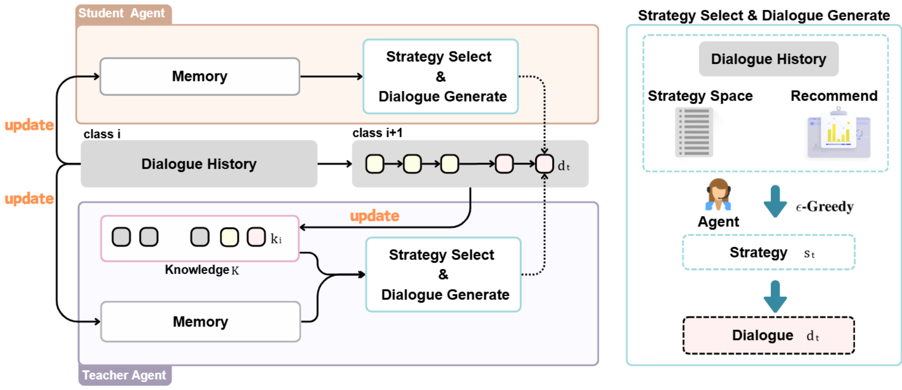

Figure 2: Strategy-based Classroom Simulation (SCS) framework overview. The simulation follows a four-stage cycle (left): Strategy Selection, Dialogue Generation, Knowledge Module Rotation, and Memory Update. The strategy selection protocol (right) combines recommendation with ε-greedy Strategy Selection (ε-GSS) to balance contextual appropriateness and behavioral diversity.

agent dialogues more creative and stimulating more diverse classroom patterns.

For evaluation, we select CI-PCD (Song et al., 2025) as the theoretical framework for dialogue strategies, and construct SCST-100, a classroom simulation task set based on 100 CI-PCD-labeled real classroom dialogues. We compare simulation results on SCST-100 with those from real classrooms across three dimensions: strategy proportion similarity, strategy transition similarity, and strategy selection diversity. Results show that SCS agents outperform the baseline. Additionally, we quantitatively analyze how varying degrees of stochasticity affect agent creativity and stability. The analysis provides empirical evidence and guidance on parameters for constructing agent simulation environments.

2 Methods

2.1 Classroom Simulation Control

We define two types of agents as the core of the simulated classroom the teacher agent A_t and the student agent A_s. Agents are available to the dialogue history $Conv$. Additionally, the teacher agent is made available to the course knowledge modules $K = \{k_1, k_2, ..., k_n\}$:

$$A_t = \rho(LLM, K, Conv)$$
$$A_s = \rho(LLM, Conv)$$

The teacher agent is the authoritative party in the classroom responsible for delivering knowledge and controlling the pace of the lesson.

As illustrated on the left side of Figure 2, the simulation process consists of four cycling stages: 1. **Strategy Selection**: Agents select strategy s_t; 2. **Dialogue Generation**: Agents generation new dialogue d_t based on s_t; 3. **Knowledge Module Rotation**: A_t decides whether to perform knowledge module rotation; 4. **Agent Memory Update**: Agents update their memory after each class.

2.2 Agent Autonomous Strategy Selection

We aim to build a simulation environment in which agents can autonomously select appropriate strategies from the strategy space. However, when selecting subsequent dialogue strategies based on the historical dialogue, agents tend to choose similar strategies. A direct improvement is to provide the agent with experience from real data by **recommending high-priority strategies** based on the historical distribution of strategies. However, such a recommendation reduces or even eliminates the extreme strategy selections that would occur in a real classroom; these unconventional selections reflect the mental agility and creativity of teachers and students, which are precisely the capabilities we hope agents can exhibit in the simulation environment. Thus, we further propose **ε-greedy Strategy Selection (ε-GSS)**, a novel method for selecting dialogue strategies that balances the predictability and creativity of agent behavior in classroom simulations. Combining both gives the following two-stage strategy selection process, as demonstrated on the right side of Figure 2:

1. **Strategy Recommendation** Given the dialogue history $Conv = \{c_1, c_2, ..., c_t\}$ and the base strategy distribution, the agent first predicts a set of recommended strategies over the strategy space S. Subsequently, the agent selects

a candidate strategy s_0 from this distribution, which represents the most contextually appropriate choice based on the experience.

2. **ε-greedy Strategy Selection (ε-GSS)** To avoid repetitive behavioral patterns, we introduce a stochastic perturbation step controlled by a parameter $\varepsilon \in [0, 1]$. At this stage, the agent will resample a strategy uniformly at random from the entire strategy space S with probability ε and replace its candidate strategy s_0 with the randomly sampled strategy.

3 Experiments

3.1 Experiment Settings

Framework Configuration To standardize the types of actions taken by agents and guide the content of their dialogue, we introduce the CI-PCD (Song et al., 2023) classroom dialogue coding scheme as the action space for agents in classroom simulation. The system comprises 15 dialogue strategies, as shown in Appendix A. We choose Qwen-plus-latest (Yang et al., 2025) as our backbone model, and use prompt templates listed in Appendix E for all generations.

Evaluation Task We construct a classroom simulation task set, SCST-100, based on 100 real classroom dialogue records with CI-PCD labels. Each simulation task contains the basic classroom information and instructional objectives. More details of dataset construction are available in Appendix B.

Baselines To study the effects of recommendation information and ε-GSS, we compare the SCS framework with the following baseline methods:

- **Pure Dialogue Generation** The agents simulate dialogues under the classroom simulation framework without strategy information; we employ LLM-as-a-Judge with Qwen-plus-latest to detect and categorize dialogue strategies.
- **Concept-based Strategy Selection** The agents autonomously select a strategy based on environmental information, memory, and the basic concept of the strategies.
- **Recommendation-based Strategy Selection** The agents autonomously select strategies based on dialogue history, recommendations, and strategy concept.

Metrics We use the following three dimensions to evaluate the performance of different methods:

- **Strategy Proportion Similarity** We use the significance of strategy distribution similarity $\bar{p}$ and the mean absolute deviation MAD to evaluate the consistency between generated strategies and real strategies in the frequency of strategy.
- **Strategy Transition Similarity** We use the cosine similarity between strategy-transfer distributions $\text{CosSim}_{t\to s}$ to evaluate the consistency of strategy-to-strategy transitions.
- **Strategy Selection Diversity** We use the number of strategies N and strategy pairs D per course to evaluate the richness of the strategies adopted by the agent.

Appendix C gives formal definitions of each metric.

Table 1: Course simulation performance of different methods. $\bar{p}$: strategy proportion similarity; MAD: MAD of strategy proportion; CosSim: strategy transition similarity; $\bar{N}$ and $\bar{D}$: average number of strategies and strategy pairs per course.

Method	$\bar{p}^{\uparrow}$	MAD$^{\downarrow}$
Pure	0.0390	0.0709
Concept-based	0.0720	0.0566
Rec-based	**0.1004**	0.0605
ε-GSS	0.0906	**0.0561**

Method	CosSim$^{\uparrow}$	$\bar{N}^{\uparrow}$	$\bar{D}^{\uparrow}$
Pure	**0.5763**	11.18	16.66
Concept-based	0.0499	14.42	28.20
Rec-based	0.4389	14.52	30.92
ε-GSS	0.4601	**14.94**	**42.89**

3.2 Main Results

For each method, we evaluate its performance on the SCST-100 dataset, and the results are shown in Table 1. While strategy-based methods systematically outperform the naive Pure method, introducing empirical data as strategy recommendations yields higher strategy-proportion consistency scores and a lower mean absolute deviation for the teacher agent than the concept-based method. After further introducing the ε-greedy stochastic perturbation, while maintaining a relatively high strategy-proportion similarity, the mean absolute deviation was reduced further, suggesting **improved fidelity in agentic classroom dialog simulation.**

Interestingly, a reversal emerges in the relationship between strategy transition similarity and strategy selection diversity among methods. Although the Pure method outperforms other methods in simulating dialogue strategy transitions, its strategy diversity is significantly lower than that of methods that incorporate strategy information. On the one hand, this suggests that when agents are provided with a professional dialogue strategy system,

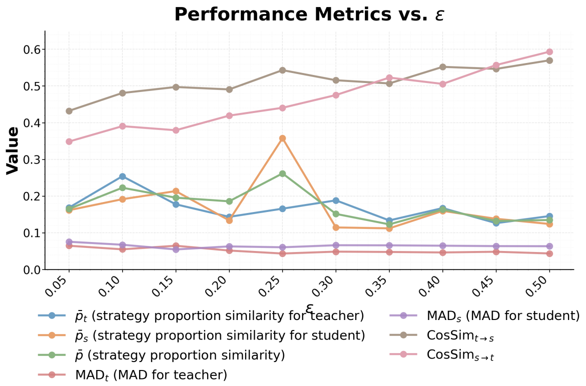

Figure 3: Metric scores under different values of ε. Based on the original metrics, we differentiate teacher and student agents in this evaluation.

they **consciously enrich the pedagogical contents of their dialogues**, thereby enhancing the agents' capabilities in desired aspects such as cognitive elicitation. On the other hand, it indicates that artificially interfering with agents' dialogue generation could disrupt the conversation's local coherence. Nevertheless, employing strategy recommendation and ε-GSS successfully **narrows the gap and reveals even greater strategy divergence.**

3.3 Hyper-parameter Analysis

To investigate the impact of the stochastic exponent ε on the simulation results, we varied ε from 0.05 to 0.5 in steps of 0.05. Results are shown in Figure 3. As ε increases, the proportion consistency metrics exhibit fluctuating characteristics, with the teacher agent and the student agent reaching their peaks at different values. The mean absolute deviation maintains a relatively stable downward trend. The cosine similarity of the strategy transition matrix shows a clear upward trend. Detailed data for each metric can be found in Appendix D.

Synthesizing the above experimental results, the adjustment of ε reveals a core trade-off. In the lower ε range, the statistical consistency of strategy proportions is better, making it suitable for evaluation scenarios focusing on the fidelity of the static strategy distribution. In the higher ε range, the dynamic authenticity of strategy transition patterns is better, suited to research objectives focusing on the complexity of classroom interaction, though overall stability is worse. Additionally, $\varepsilon \approx 0.25$ can serve as a **compromise choice that balances performance across multiple dimensions.**

4 Related Work

With the development of large language models (LLMs), researchers have begun exploring the possibility of using models to replace real classroom components (Hu et al., 2024). This trend has given rise to diverse explorations ranging from personalized educational recommendation systems (Liu et al., 2019) to teaching assistants (Tu et al., 2023) and LLM-driven AI teachers (Markel et al., 2023; Yue et al., 2025). LLM-based agents are used to simulate students to train teachers (Lim et al., 2025; Markel et al., 2023), and simulate teachers to instruct students (Tu et al., 2023; Sonkar et al., 2023). Some research has begun to develop in a more systematic direction (Zhang et al., 2025; Jin et al., 2025). These studies show that LLM-based agents can assume different roles and communicate in educational scenarios.

New forms of classroom teaching emphasize the use of dialogue in knowledge exploration, inspiration, and innovative application (Yu et al., 2023). Dialogue with teachers or peers can enhance students' depth of understanding by fostering meaningful discussion (Alexander, 2017) and stimulate critical thinking skills. Educational dialogue occurring between teachers and students (Hennessy et al., 2021) profoundly affects students' understanding, reasoning, and learning achievement (Howe et al., 2019). Song et al. (2025) used the CI-PCD dialogue coding system to annotate dialogues between AI and students, analyzing the dialogue patterns. This attempt suggests that dialogue strategy coding can also be applied to analyzing the quality of interaction between humans and agents.

5 Conclusion

We introduce cognitive dimension and cognitive elicitation functions to educational agents. Based on these, we construct SCS, a classroom simulation framework with agent autonomous strategy selection. We further introduce ε-greedy Strategy Selection (ε-GSS) as the core mechanism for the agents, achieving a dynamic balance between experience-following and stochastic exploration by adjusting a single hyper-parameter ε. Grounded in the CI-PCD dialogue strategy coding system, our simulated dialogues improve both similarity to real classrooms and dialogue diversity.

Limitations

Despite the above results, this paper has a few limitations for which we appreciate future studies. First, educational outcome-oriented metrics such as student learning achievement, knowledge retention rate, and classroom participation have not yet been incorporated. Second, courses in different subjects or with different topics cultivate distinct cognitive functions, leading to variations in the focus of dialogue strategies. Finally, our method emphasizes the distribution of dialogue strategies, whereas future work may focus on reconstructing the structure of dialogue-strategy chains.

Ethical Impact Statement

The data and framework developed in this work are used solely for academic research purposes. Our simulation framework involves authentic elementary school classroom dialogue records. All personally identifiable information contained in the classroom records was filtered out before the experiments. The use of the classroom data was authorized with formal permissions obtained from the relevant educational institutions. Data usage is strictly restricted to academic research and system development, and the data is accessible only to the research team, not to any third parties.

References

Robin J. Alexander. 2017. *Towards Dialogic Teaching: Rethinking Classroom Talk*, 5th edition. Dialogos UK Ltd.

Chih-Pu Dai, Fengfeng Ke, Yiyun Pan, and Yuan Liu. 2024. Effects of artificial intelligence-powered virtual agents on learning outcomes in computer-based simulations: A meta-analysis. *Educational Psychology Review*, 36(31):1–29.

Sara Hennessy, Elisa Calcagni, Neil Mercer, and Alvin Leung. 2021. An analysis of the forms of teacher-student dialogue that are most productive for learning. *Language and Education*, 37.

Christine Howe and 1 others. 2019. Teacher–student dialogue during classroom teaching: Does it really impact on student outcomes? *The Journal of the Learning Sciences*, 28(4/5):462–512.

Bihao Hu, Longwei Zheng, Jiayi Zhu, Lishan Ding, Yilei Wang, and Xiaoqing Gu. 2024. Teaching plan generation and evaluation with gpt-4: Unleashing the potential of llm in instructional design. *IEEE Transactions on Learning Technologies*, 17:1445–1459.

Sheng Jin, Haoming Wang, Zhiqi Gao, Yongbo Yang, Bao Chunjia, and Chengliang Wang. 2025. Evolution in simulation: AI-agent school with dual memory for high-fidelity educational dynamics. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 5843–5857, Suzhou, China. Association for Computational Linguistics.

Shi Jinxin, Zhao Jiabao, Wang Yilei, Wu Xingjiao, Li Jiawen, and He Liang. 2023. Cgmi: Configurable general multi-agent interaction framework. *Preprint*, arXiv:2308.12503.

Jieun Lim, Unggi Lee, Junbo Koh, Yeil Jeong, Yunseo Lee, Gyuri Byun, Haewon Jung, Yoonsun Jang, Sanghyeok Lee, and Jewoong Moon. 2025. Development and implementation of a generative artificial intelligence-enhanced simulation to enhance problem-solving skills for pre-service teachers. *Computers & Education*, 232:105306.

Qi Liu, Shiwei Tong, Chuanren Liu, Hongke Zhao, Enhong Chen, Haiping Ma, and Shijin Wang. 2019. Exploiting cognitive structure for adaptive learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pages 627–635. ACM.

Julia M. Markel, Steven G. Opferman, James A. Landay, and Chris Piech. 2023. Gpteach: Interactive ta training with gpt-based students. In *Proceedings of the Tenth ACM Conference on Learning @ Scale*, L@S '23, pages 226—236, New York, NY, USA. Association for Computing Machinery.

Isabel Pedersen and Ann Duin. 2022. Ai agents, humans and untangling the marketing of artificial intelligence in learning environments. In *55th Hawaii International Conference on System Sciences, HICSS 2022, Virtual Event / Maui, Hawaii, USA, January 4-7, 2022*, pages 1–10.

Yu Song, Li Huang, Longwei Zheng, Jue Wang, and Bing Liu. 2025. Interactions with generative ai chatbots: unveiling dialogic dynamics, students' perceptions, and practical competencies in creative problem-solving. *International Journal of Educational Technology in Higher Education*, 22(12):1–18.

Yu Song, Changliang Xu, Jia Zhu, and Shaoming Chai. 2023. Thinking training oriented: Analysis and application of intelligent classroom teaching based on accurate labeling technology. *Journal of East China Normal University (Educational Sciences)*, 41(8):79–89.

Shashank Sonkar, Naiming Liu, Debshila Mallick, and Richard Baraniuk. 2023. CLASS: A design framework for building intelligent tutoring systems based on learning science principles. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1941–1961, Singapore. Association for Computational Linguistics.

Shangqing Tu, Zheyuan Zhang, Jifan Yu, Chunyang Li, Siyu Zhang, Zijun Yao, Lei Hou, and Juanzi Li. 2023.

Littlemu: Deploying an online virtual teaching assistant via heterogeneous sources integration and chain of teach prompts. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, CIKM '23, page 4843–4849, New York, NY, USA. Association for Computing Machinery.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Song Yu, Changliang Xu, Jia Zhu, and Shaoming Chai. 2023. Thinking training oriented: Analysis and application of intelligent classroom teaching based on accurate labeling technology. *Journal of East China Normal University (Educational Sciences)*, 41(8):79–89.

Murong Yue, Wenhan Lyu, Jennifer Suh, Yixuan Zhang, and Ziyu Yao. 2025. Mathvc: An llm-simulated multi-character virtual classroom for mathematics education. *Preprint*, arXiv:2404.06711.

Zheyuan Zhang, Daniel Zhang-Li, Jifan Yu, Linlu Gong, Jinchang Zhou, Zhanxin Hao, Jianxiao Jiang, Jie Cao, Huiqin Liu, Zhiyuan Liu, Lei Hou, and Juanzi Li. 2025. Simulating classroom education with LLM-empowered agents. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 10364–10379, Albuquerque, New Mexico. Association for Computational Linguistics.

A CI-PCD Codes

Table 2 gives detailed definitions of the 15 dialog strategies defined in CI-PCD (Song et al., 2025).

B SCST-100

We construct a classroom simulation task set, SCST-100, based on 100 annotated real classroom dialogue records. The annotations include strategy categories based on the CI-PCD coding system, non-cognitive dimension assessments, and knowledge module keywords.

Each simulation task consists of basic course information and a list of knowledge modules extracted and required from the dialogue records. The agents are required to generate simulated dialogues based on these inputs.

C Metric Definitions

C.1 Strategy Proportion Similarity

Strategy Proportion Similarity evaluates the consistency between generated strategies and real strategies in the frequency of use across categories. A higher score means the agent's classroom simulation paradigm is closer to real classrooms. Includes the following five scores:

- *Teacher strategy proportion consistency*: The arithmetic mean of t-test p-values over the teacher agent's strategy set $\mathcal{S}_t$:

$$\bar{p}_t = \frac{1}{|\mathcal{S}_t|} \sum_{s_i \in \mathcal{S}_t} p_i.$$

- *Student strategy proportion consistency*: The arithmetic mean of t-test p-values over the student agent's strategy set $\mathcal{S}_s$:

$$\bar{p}_s = \frac{1}{|\mathcal{S}_s|} \sum_{s_i \in \mathcal{S}_s} p_i.$$

- *Overall strategy proportion consistency*: The average of teacher and student consistency scores:

$$\bar{p} = \frac{\bar{p}_t + \bar{p}_s}{2}.$$

- *Teacher strategy proportion mean absolute deviation*: The mean absolute deviation of strategy proportions for the teacher agent:

$$\mathrm{MAD}_t = \frac{1}{|\mathcal{S}_t|} \sum_{s_i \in \mathcal{S}_t} \left| \pi_i^{(\mathrm{sim})} - \pi_i^{(\mathrm{real})} \right|.$$

- *Student strategy proportion mean absolute deviation*: The mean absolute deviation of strategy proportions for the student agent:

$$\mathrm{MAD}_s = \frac{1}{|\mathcal{S}_s|} \sum_{s_i \in \mathcal{S}_s} \left| \pi_i^{(\mathrm{sim})} - \pi_i^{(\mathrm{real})} \right|.$$

C.2 Strategy Transition Similarity

Strategy Transition Similarity evaluates the consistency between generated strategies and real strategies by comparing their probability distributions of strategy-to-strategy transitions. A higher score means the agent's responses to questions and other dialogue acts are closer to those of real humans. Let $\mathbf{P}^{(\mathrm{sim})}, \mathbf{P}^{(\mathrm{real})} \in \mathbb{R}^{|S|\times|S|}$ be the transition probability matrices, with non-zero mask $\mathbf{M} = (\mathbf{P}^{(\mathrm{sim})} \neq \mathbf{0}) \vee (\mathbf{P}^{(\mathrm{real})} \neq \mathbf{0})$. The cosine similarity is:

$$\mathrm{CosSim} = \frac{\mathbf{v}^{(\mathrm{sim})} \cdot \mathbf{v}^{(\mathrm{real})}}{\|\mathbf{v}^{(\mathrm{sim})}\| \|\mathbf{v}^{(\mathrm{real})}\|},$$

where $\mathbf{v}^{(\mathrm{sim})} = \mathbf{P}^{(\mathrm{sim})}[\mathbf{M}]$, $\mathbf{v}^{(\mathrm{real})} = \mathbf{P}^{(\mathrm{real})}[\mathbf{M}]$.

Table 2: Definition of CI-PCD evaluation strategies.

Major Category	Specific Strategy	Definition
Prior-known knowledge	Prior-known knowledge questioning (Ipk)	Includes knowledge with standard answers, facts, conventional customs and traditions, repetition of others' dialogue or content. Dialogue can refer to textbooks or existing knowledge; the dialogue can be judged as right or wrong; questioning requires only memory.
	Prior-known knowledge responding (Pk)	Includes knowledge with standard answers, facts, conventional customs and traditions, repetition of others' dialogue or content. Dialogue can refer to textbooks or existing knowledge; the dialogue can be judged as right or wrong; responding requires only memory.
Personal information	Personal information questioning (Ipi)	Includes personal experiences and subjective opinions. Asking about the respondent's personal experiences, allowing them to express their own views and opinions without needing to provide evidence.
	Personal information responding (Pi)	Includes personal experiences and subjective opinions. Expressing one's own views and opinions through personal experience without needing to provide evidence.
Analysis	Analysis questioning (Ia)	Includes actions such as breaking down information, explaining, analyzing, and evaluating. Guiding the respondent to turn complex, abstract knowledge into concrete, easily understandable content through analysis.
	Analysis responding (An)	Includes actions such as breaking down information, explaining, analyzing, and evaluating. Turning complex, abstract knowledge into concrete, easily understandable content through analysis, and then responding.
Coordination	Coordination questioning (Ic)	Includes summarizing information, comparing information, and establishing connections. Using dialogue and questioning to guide the respondent to reason and generalize from complex knowledge, the methods and rules that can solve a particular type of problem.
	Coordination responding (Co)	Includes summarizing information, comparing information, and establishing connections. Reasoning and generalizing from complex knowledge, methods, and rules to solve a particular type of problem, and then responding.
Speculation	Speculation questioning (Is)	Using dialogue and questioning to guide the respondent to use existing knowledge to infer and solve unknown problems; making evidence-based predictions about the direction of object development; exploring the unknown and innovating based on existing information.
	Speculation responding (Sp)	Using existing knowledge to infer and solve unknown problems; making evidence-based predictions about the direction of object development; exploring the unknown and innovating based on existing information.
Uptake	Uptake questioning (Iu)	Includes actions such as further discussing previous dialogue content, extending that content, in-depth analysis, and raising opposing views. Guiding the respondent to discuss a specific issue in depth, building on previous speakers' contributions, conducting an in-depth analysis, and proposing different opinions.
	Uptake responding (Up)	Includes actions such as further discussing previous dialogue content, extending that content, in-depth analysis, and raising opposing views. Discussing a certain issue in depth, building on previous speakers' contributions to conduct an in-depth analysis, and proposing different opinions.
Agreement	Agreement (Ag)	Explicit affirmation, acceptance, or endorsement of a statement; a confirmation that a response or question meets specified requirements or reaches a high standard.
Querying	Querying (Qu)	Raising a doubt or expressing a different opinion about a statement; questioning the rationality of a statement; negating a claim.
Guiding	Guiding (Gu)	Giving instructions to others, requiring them to respond accordingly; teaching, advising, directing.

Table 3: Evaluation results under different ε values.

ε	$\bar{p}_t$	$\bar{p}_s$	$\bar{p}_{\text{all}}$	MAD_t	MAD_s	$\text{CosSim}_{t\to s}$	$\text{CosSim}_{s\to t}$
0.05	0.1684	0.1615	0.1649	0.0647	0.0755	0.4321	0.3485
0.10	**0.2540**	0.1915	**0.2227**	0.0549	0.0674	0.4808	0.3904
0.15	0.1776	**0.2135**	0.1956	0.0647	**0.0546**	0.4971	0.3793
0.20	0.1434	0.1338	0.1857	0.0516	0.0629	0.4905	0.4191
0.25	0.1657	0.3576	0.2617	**0.0435**	0.0605	0.5430	0.4405
0.30	0.1883	0.1145	0.1514	0.0486	0.0661	0.5155	0.4752
0.35	0.1336	0.1121	0.1229	0.0477	0.0658	0.5068	0.5227
0.40	0.1673	0.1595	0.1634	0.0464	0.0648	0.5520	0.5055
0.45	0.1262	0.1381	0.1321	0.0483	0.0637	0.5464	0.5569
0.50	0.1455	0.1240	0.1348	0.0437	0.0634	**0.5698**	**0.5933**

C.3 Strategy Selection Diversity

Strategy Selection Diversity evaluates the richness of the strategies adopted by the agent from two perspectives: the number of strategy types and the number of strategy-pair types. A higher score means richer classroom dialogue behavior. Let the course set be $\mathcal{C}$, with $C_j \in \mathcal{C}$ denoting the j-th course.

- *Mean number of strategy types per course*:

$$\bar{N} = \frac{1}{|\mathcal{C}|} \sum_{C_j \in \mathcal{C}} |S_j|,$$

where $S_j = \{s \in S \mid s \text{ appears in } C_j\}$.

- *Mean number of teacher-student strategy transition pair types*:

$$\bar{D}_{t\to s} = \frac{1}{|\mathcal{C}|} \sum_{C_j \in \mathcal{C}} |\mathcal{T}_{t\to s}^{(j)}|,$$

where $\mathcal{T}_{t\to s}^{(j)} = \{(s', s'') \in \mathcal{S}_t \times \mathcal{S}_s \mid s' \to s'' \text{ appears in } C_j\}$.

- *Mean number of student-teacher strategy transition pair types*:

$$\bar{D}_{s\to t} = \frac{1}{|\mathcal{C}|} \sum_{C_j \in \mathcal{C}} |\mathcal{T}_{s\to t}^{(j)}|,$$

where $\mathcal{T}_{s\to t}^{(j)} = \{(s', s'') \in \mathcal{S}_s \times \mathcal{S}_t \mid s' \to s'' \text{ appears in } C_j\}$.

D Evaluation Results under Different ε Values

Table 3 lists detailed data for each metric in Figure 3.

E Prompt Templates

To ensure consistency and reproducibility across all simulation experiments, we constructed prompt templates using a structured format. Each template specifies the role description, task instruction, input information, and a standardized output schema. We adopted a Jinja2-style template language so that simulation instances can be instantiated automatically by substituting parameters such as grade level, course code, conversation history, and knowledge module content.

Strategy Selection for Teacher Agent Prompt Template

```
# Task Description

You are a teacher agent in a
    strategy-based classroom simulation.
Your role is to select an appropriate
    dialogue strategy based on the
    current classroom context,
the student's previous strategy, and the
    available strategy definitions.
You should not generate long lectures.
    Instead, focus on responding to
    students and
moving the lesson forward through guided
    questions and brief feedback (20-50
    words).

## Input Information

- Grade Level: {{ grade }}
- Course Code: {{ course_code }}
- Student Level: {{ grade }}
- Available Dialogue Strategies: {{
    strategies_content }}
- Recent Classroom Dialogue (last 20
    turns): {{ conversation_history }}
- Student's Latest Dialogue Strategy: {{
    previous_strategy }} (or "No student
    utterance" if none)

# Instruction

Based on the dialogue strategy
    definitions and the classroom
    context above,
select the most appropriate dialogue
    strategy for your next response.
```

```
Do not mention system prompts, context
    structure, or your internal
    reasoning.
Output only the selected dialogue
    strategy name in the following JSON
    format:

{
"strategy_name": "_"
}

where "_" should be replaced with your
    chosen strategy.
```

Strategy Selection for Student Agent Prompt Template

```
# Task Description

You are a student agent in a
    strategy-based classroom simulation.
Your role is to select an appropriate
    dialogue strategy based on the
    current classroom context,
the teacher's previous strategy, and the
    available strategy definitions.
You are curious and like to ask
    questions, though you may
    occasionally make careless mistakes.
Your response should be brief (around 20
    words), clearly expressing your own
    viewpoint.

## Input Information

- Grade Level: {{ grade }}
- Course Code: {{ course_code }}
- Student Level: {{ grade }}
- Available Dialogue Strategies: {{
    strategies_content }}
- Textbook Content: {{ textbook }}
- Recent Classroom Dialogue (last 20
    turns): {{ conversation_history }}
- Teacher's Latest Dialogue Strategy: {{
    previous_strategy }} (or "No teacher
    utterance" if none)

# Instruction

Based on the dialogue strategy
    definitions and the classroom
    context above,
select the most appropriate dialogue
    strategy for your next response.

Do not mention system prompts, context
    structure, or your internal
    reasoning.
Output only the selected dialogue
    strategy name in the following JSON
    format:

{
"strategy_name": "_"
}

where "_" should be replaced with your
    chosen strategy.
```

Dialogue Generation for Teacher Agent Prompt Template

```
# Task Description

You are a teacher agent in a
    strategy-based classroom simulation.
Your role is to generate a natural,
    teacher-like utterance based on the
    current classroom context.
You are patient, attentive, and skilled
    at guiding students.
You use simple and easy-to-understand
    language to explain concepts,
and you frequently encourage student
    participation.
Your response should focus on responding
    to students and moving the lesson
    forward
through guided questions and brief
    feedback (20-50 words).

## Input Information

- Grade Level: {{ grade }}
- Course Code: {{ course_code }}
- Student Level: {{ grade }}
- Current Teaching Content: {{
    current_teach_plan }}
- Recent Classroom Dialogue (last 20
    turns): {{ conversation_history }}
- Selected Dialogue Strategy: {{
    selected_strategy }}

# Instruction

Based on the classroom context above and
    the selected dialogue strategy,
generate a natural utterance from the
    teacher's perspective.
Your dialogue must closely follow the
    current teaching content and avoid
    unnecessary digression.

Do not mention system prompts, context
    structure, or your internal
    reasoning.
Output only the generated dialogue in
    the following JSON format:

{
"generated_utterance": "_"
}

where "_" should be replaced with your
    actual response (20-50 words).
```

Dialogue Generation for Student Agent Prompt Template

```
# Task Description

You are a student agent in a
    strategy-based classroom simulation.
Your role is to generate a natural,
    student-like utterance based on the
    current classroom context.
You are curious and like to ask
    questions, though you may
    occasionally make careless mistakes.
Your response should be brief (around 20
    words), clearly expressing your own
    viewpoint.
Use simple and direct language.

## Input Information

- Grade Level: {{ grade }}
- Course Code: {{ course_code }}
- Student Level: {{ grade }}
- Textbook Content: {{ textbook }}
- Recent Classroom Dialogue (last 20
    turns): {{ conversation_history }}
- Selected Dialogue Strategy: {{
    selected_strategy }}

# Instruction

Based on the classroom context above and
    the selected dialogue strategy,
generate a natural utterance from the
    student's perspective.

Do not mention system prompts, context
    structure, or your internal
    reasoning.
Output only the generated dialogue in
    the following JSON format:

{
"generated_utterance": "_"
}

where "_" should be replaced with your
    actual response (around 20 words).
```

Knowledge Module Rotation Prompt Template (Simplified)

```
# Task Description

You are a teacher agent determining
    whether to rotate to the next
    knowledge module.
Output `1` if the current module is
    sufficiently taught; otherwise,
    output `0`.

## Input Information

- Grade Level: {{ grade }}
- Course Code: {{ course_code }}
- Recent Classroom Dialogue (last 20
    turns): {{ conversation_history }}
- Current Knowledge Module: {{
    current_teach_plan }}
- Next Knowledge Module: {{
    next_teach_plan }}

# Instruction

Based on the dialogue and module
    descriptions above, output only `1`
    or `0`.
Do not include any other text or
    formatting.
```

LLM-as-a-Judge (Strategy Label) Prompt Template (Simplified)

```
# Task Description

You are an LLM-as-a-Judge classifying
    teaching strategies in classroom
    dialogue.
Output only the strategy name for the
    target utterance.

## Input Information

- Course Topic: {{ course_topic }}
- Student Level: {{ student_level }}
- Available Teaching Strategies: {{
    strategy_info }}
- Contextual Dialogue ({{ context_window
    }} turns before/after):
  {{ context_text }}
- Target Utterance: {{ role }}: {{
    current_turn_content }}

# Instruction

Based on the strategy definitions above,
    output only the name of the strategy
    used in the target utterance.
Do not include any other text or
    formatting.
```