

# Interpretable Difficulty-Aware Knowledge Tracing in Tutor-Student Dialogues

Shuyan Huang Alexander Scarlatos Jaewook Lee Andrew Lan

University of Massachusetts Amherst

{shuang, ajscarlatos, jaewooklee, andrewlan}@umass.edu

## Abstract

Recent advances in large language models (LLMs) have led to the development of AI-powered tutoring systems that provide interactive support via dialogue. To enable these tutoring systems to provide personalized support, it is essential to assess student performance at each turn, motivating knowledge tracing (KT) in dialogue settings. However, existing dialogue-based KT approaches often ignore question difficulty and rely on opaque LLM latent representations, hindering accurate and interpretable prediction. In this work, we propose an interpretable difficulty-aware conversational KT framework that leverages LLMs to explicitly model student knowledge state and the difficulty of tutor-posed tasks at each dialogue turn. The framework incorporates the original question and the next tutor-posed task to estimate the student's knowledge state and the difficulty of the upcoming turn. It further integrates Item Response Theory to map LLM outputs into student ability and question difficulty parameters, enabling interpretable prediction of student performance grounded in cognitive theories of learning. We evaluate the framework on two tutor-student dialogue datasets. Quantitative and qualitative results show that our framework outperforms existing KT baselines, meanwhile generating interpretable outputs consistent with cognitive theory<sup>1</sup>.

## 1 Introduction

Tutoring, which provides tailored instruction and feedback to students in an interactive manner, has been shown to effectively improve learning outcomes (Nickow et al., 2020; Nye et al., 2014). With the advanced natural language understanding and generation capabilities of large language models (LLMs), researchers have developed an increasing number of LLM-based tutoring systems, such as

<sup>1</sup>Our code and data are available at <https://github.com/umass-ml4ed/Difficulty-Aware-DialogKT>

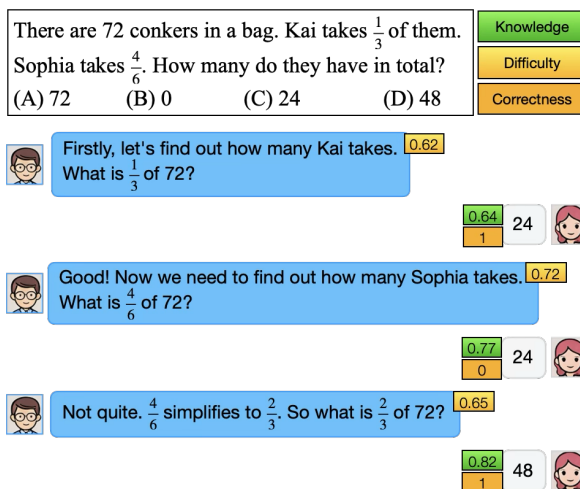


Figure 1: An example dialogue illustrating a student's learning trajectory as the tutor adapts difficulty while scaffolding the math problem.

Khanmigo (Khan Academy, 2023) and LiveHint (Carnegie Learning, 2024). Recent work has further explored how to improve LLM-based tutors by training them to optimize student learning outcomes and by steering them toward diverse tutor personas that reflect different scaffolding, feedback, and affective-support styles (Scarlatos et al., 2025c; Lee et al., 2026). These systems can significantly reduce the workload of human tutors and provide scalable support for a large number of students.

To provide personalized pedagogical support, it is crucial to assess students' mastery of concepts based on their responses throughout the dialogue. The process of modeling students' knowledge states over time to predict their future performance is known as knowledge tracing (KT). Recently, with the development of AI-powered tutoring systems, KT has been extended to dialogue-based settings to predict whether students can correctly answer subsequent tutor-posed tasks (Scarlatos et al., 2025a; Wang et al., 2025).

However, existing KT approaches in dialogue

settings have two important limitations. First, prior studies do not explicitly model question difficulty. In practice, the correctness of student responses is not solely determined by their knowledge state but is also substantially influenced by question difficulty (Shen et al., 2022). Even students with high knowledge mastery may fail to answer very difficult questions. For example, as shown in Figure 1, although the student has a relatively high estimated knowledge state (e.g., 0.64), the student still answers the next tutor-posed task incorrectly when the question difficulty increases (e.g., from 0.62 to 0.72). Prior studies that model student responses solely based on students’ knowledge states may yield inaccurate predictions (Scarlatos et al., 2025a). Second, the reasoning process underlying model predictions of students’ future performance remains opaque. Most current approaches rely on LLMs with high-dimensional latent representations that lack interpretability. Consequently, it is difficult to derive psychologically meaningful explanations that align with cognitive theories of learning. This lack of transparency may undermine tutors’ trust in model predictions, hindering the adoption of LLMs as effective tutoring tools. As illustrated in Figure 1, even when the student answers a more difficult task incorrectly, the estimated knowledge state continues to increase, reflecting gradual learning dynamics consistent with cognitive theories (Newell and Rosenbloom, 1981; Anderson, 2000).

### 1.1 Contributions

To address the aforementioned challenges, we propose an interpretable difficulty-aware conversational KT framework built upon an LLM-based architecture that explicitly models both student ability and tutor-turn difficulty. Specifically, we prompt an LLM with the question and dialogue context to estimate student knowledge, and further prompt the model with the tutor turn to estimate the difficulty of the tutor-posed task. Furthermore, to bridge the gap between model predictions and cognitive theories of learning, we integrate principles from Item Response Theory (IRT). More specifically, we feed the estimated ability and difficulty into an IRT model to predict student correctness, enabling fine-tuning of the LLM using observed correctness labels. This design aligns model predictions with cognitive theory and improves interpretability by providing transparent reasoning based on the relationship between student ability and task difficulty. To evaluate the effectiveness of our approach, we

conduct comprehensive experiments on two tutor-student dialogue datasets. Both quantitative and qualitative results demonstrate the superiority of our framework in student ability and difficulty estimation throughout tutoring dialogues.

## 2 Problem Statement

Practically, a dialogue often starts when a student requests assistance from a tutor when they cannot successfully solve a math question/problem. Formally, we represent each dialogue  $\mathcal{D}$  as an ordered sequence of alternating tutor turns  $t$  and student turns  $s$ , i.e.,  $\mathcal{D} = (s_0, t_1, s_1, \dots, t_N, s_N)$ , where  $N$  denotes the total number of turn pairs in the dialogue and  $s_0$  is included when the student initiates the dialogue. In cases where the same speaker produces several consecutive utterances, we merge them into a single tutor turn or student turn to maintain the alternating turn structure. Following Scarlatos et al. (2025a), each dialogue can be treated as a sequence of formative assessments, where each tutor turn  $t_j$  (for turn index  $j$ ) acts as an assessment question, and the corresponding student turn  $s_j$  represents a correct or incorrect response. This perspective allows us to use tutor-student dialogues to model student learning over time. Each tutor turn with pedagogical intent (i.e., designed to assess student knowledge) is associated with one or more related knowledge components (KCs),  $\mathcal{C}_j = \{c_j^k\}_{k=1}^{K_j}$ , where  $K_j$  is the number of KCs associated with turn  $t_j$ . For student turns that follow these tutor turns, we define the correctness of a student turn  $s_j$  as  $r_j \in \{0, 1\}$ . In this way, we formulate the dialogue-based KT task as predicting the correctness  $\hat{r}_{j+1}$  of the student’s next response with the given question  $q$ , historical dialogue  $(t_1, s_1, \dots, t_j, s_j)$ , and next tutor-posed task  $t_{j+1}$ , which can be formally defined as:

$$\hat{r}_{j+1} \sim f_{\Theta}(q, t_1, s_1, \dots, t_j, s_j, t_{j+1}), \quad (1)$$

where  $\Theta$  denotes the learnable parameters of the framework, which are optimized to predict students’ future performance.

## 3 Approach

In this section, we detail our framework that consists of three modules: (1) a *knowledge estimator* that extracts latent knowledge representations from both the textual content of the question and the tutor-student dialogue history; (2) a *difficulty*

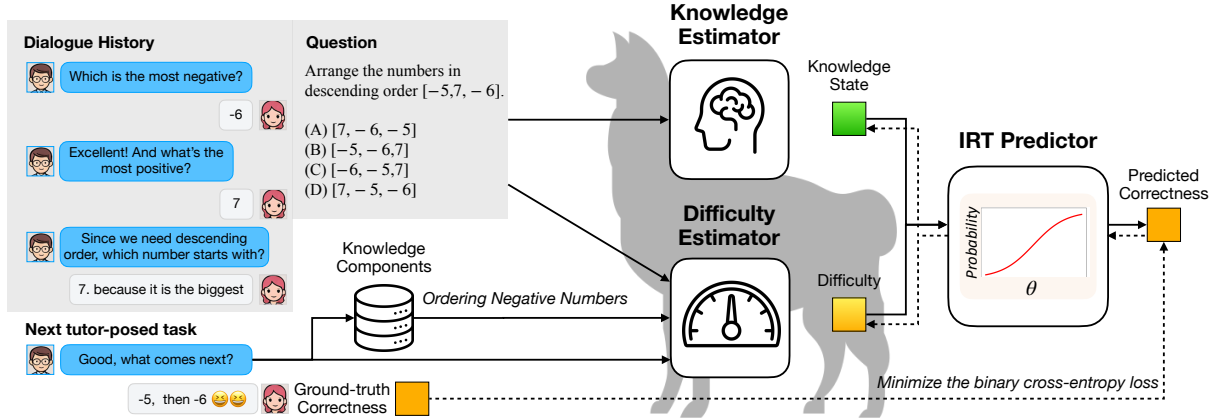


Figure 2: Overview of our framework with three modules: a knowledge estimator, a difficulty estimator, and an IRT-based predictor that jointly model student knowledge and task difficulty to predict correctness on the next task.

*estimator* that computes the difficulty of the tutor-posed task based on the question content, the dialogue history and content of the current turn; and (3) an *IRT-based predictor* that integrates the estimated knowledge state and difficulty representation to provide interpretable predictions of whether the student will correctly respond to the tutor-posed task. We illustrate our full framework in Figure 2, and show our prompts in Appendix A.

### 3.1 Knowledge Estimator

We now introduce the knowledge estimator module of our framework, which estimates student knowledge states for each turn in a dialogue. At each turn, the student’s subsequent performance depends on both the original question being solved and the dialogue history up to the current turn. The original question provides necessary context, while the dialogue history reveals the student’s evolving knowledge state, including observed errors or misconceptions. To effectively represent the textual information, we leverage the natural language understanding capabilities of LLMs to jointly encode the question text and dialogue history. To adapt general-purpose LLMs to the educational domain, we further fine-tune the model on dialogue-based tutoring data. We construct a prompt  $I_{j+1}^K$  representing turn  $j + 1$ , where we instruct the LLM to generate a one-word prediction for the student’s current ability level. Prior LLM-based KT work categorizes students’ knowledge states into discrete levels, e.g., good/fail (Li et al., 2025b). Inspired by this classification design, we use “GOOD” and “BAD” as positive and negative verbalizer tokens following prompt-based classification methods (Schick and Schütze, 2021). Rather than sampling text from the

LLM, we construct a continuous ability estimate by comparing their next-token logits. We first extract the LLM’s next-token output logits corresponding to these two verbalizer tokens:

$$z_{j+1}^{\text{GOOD}}, z_{j+1}^{\text{BAD}} = \text{LLM}(I_{j+1}^K).$$

We convert these logits into a scalar ability estimate by taking their difference, defining the student’s ability at the next turn as:

$$\theta_{j+1} = z_{j+1}^{\text{GOOD}} - z_{j+1}^{\text{BAD}},$$

where  $\theta_{j+1}$  represents the student’s latent knowledge state at turn  $j + 1$ .

### 3.2 Difficulty Estimator

We now introduce the difficulty estimator module of our framework, which estimates the difficulty of each tutor-posed task in a dialogue. Difficulty estimation is recognized as a fundamental component in classical assessment settings, where accurately modeling item difficulty improves the reliability of student evaluation (Scarlatos et al., 2025b). In dialogue-based tutoring, each interaction is typically centered around an original question that the student is attempting to solve with tutor assistance. This original question is associated with an inherent difficulty determined by its textual content and cognitive complexity. As the dialogue progresses, the tutor can actively pose tasks with different difficulty in subsequent turns. Hence, to more accurately predict student performance at each turn, we estimate the difficulty of each tutor-posed task by jointly considering the original question text, dialogue context, and the next tutor-posed task. Similar to the knowledge estimator, we construct a prompt  $I_{j+1}^D$

representing turn  $j+1$ , where we instruct the model to generate a one-word prediction for the difficulty level of the tutor-posed task. Again, we extract the LLM’s output logits, now corresponding to the “HARD” and “EASY” verbalizer tokens:

$$z_{j+1}^{\text{HARD}}, z_{j+1}^{\text{EASY}} = \text{LLM}(I_{j+1}^{\text{D}}).$$

We convert the logits into a scalar difficulty estimate, defining the difficulty of the tutor-posed task:

$$d_{j+1} = z_{j+1}^{\text{HARD}} - z_{j+1}^{\text{EASY}},$$

where  $d_{j+1}$  represents the latent difficulty of the tutor-posed task at turn  $j+1$ .

### 3.3 IRT-based Predictor

In prior LLM-based dialogue KT approaches, it is very challenging to explain the LLMs’ parameters and their decision-making processes when predicting student response correctness. Therefore, to align LLMs’ predictions with models in psychometric measurement, we introduce an IRT-based prediction layer that jointly predicts student response correctness while providing interpretable estimates of student knowledge and tutor turn difficulty. IRT, particularly the one-parameter logistic (1PL) model, also known as the Rasch model (Rasch, 1993), is widely used in educational assessment to model the probability of a correct response based on student ability and task difficulty. In the 1PL IRT model, student ability and task difficulty are represented as latent variables, and the probability of correctness increases as the positive difference between ability and difficulty grows. This formulation provides interpretability since both student knowledge and task difficulty are explicitly represented as meaningful cognitive quantities. Given the estimated ability  $\theta_{j+1}$  and task difficulty  $d_{j+1}$ , the probability that a student will answer correctly at turn  $j+1$  is defined as:

$$\hat{r}_{j+1} = \frac{1}{1 + \exp(-\alpha(\theta_{j+1} - d_{j+1}))},$$

where  $\alpha$  is a learnable scalar that controls the sensitivity of the probability to differences between ability and difficulty. This parameter plays a role similar to the discrimination parameter in the two-parameter logistic (2PL) IRT model (Birnbaum, 1968), enabling the model to adjust how strongly performance depends on the gap between ability and difficulty. Please note that we additionally experimented with modeling  $\alpha$  as a learnable vector

to capture KC-specific discrimination, analogous to the 2PL model. However, this setting led to degraded performance in our experiments, likely due to overparameterization. Therefore, we adopt a scalar  $\alpha$  in our final framework.

#### 3.3.1 Model Optimization

To improve predictive accuracy, we fine-tune the LLM on correctness in observed tutor-student dialogues. We optimize the learnable parameters of the model by minimizing the binary cross-entropy loss between the ground-truth correctness of student responses  $r_{j+1}$  and the predicted probabilities  $\hat{r}_{j+1}$  produced by the IRT-based prediction layer:

$$\mathcal{L} = -\sum (r_{j+1} \log \hat{r}_{j+1} + (1 - r_{j+1}) \log (1 - \hat{r}_{j+1})).$$

## 4 Experiments

### 4.1 Datasets

To evaluate the effectiveness of our framework, we conduct experiments on two datasets as follows:

**Question-Anchored Tutoring Dialogues (QATD<sub>2k</sub>)** QATD<sub>2k</sub> is the largest open-source real-world tutoring dialogue dataset (Zent et al., 2025). The dataset is collected from one-on-one conversations between tutors and students on the Eedi mathematics learning platform. Each dialogue begins when a student requests assistance in solving a math multiple-choice problem. After data preprocessing, the dataset contains 1,573 and 393 dialogues in the training set and test set, respectively. We use the LLM-generated KC and response correctness annotations collected in prior work (Scarlatos et al., 2026).

**MathDial** MathDial is a tutoring dataset grounded in multi-step math reasoning problems between GPT-3.5-simulated students and crowd workers acting as tutors (Macina et al., 2025). Each dialogue begins with an incorrect student solution and the tutor guides the student to identify and correct their errors. The dataset contains 2,235 and 588 dialogues in the training set and test set, respectively. We use the LLM-generated KC and response correctness annotations collected in prior work (Scarlatos et al., 2025a).

### 4.2 Baselines

We compare our framework against two types of representative KT baselines: LLM-based models and deep learning-based models.

### 4.2.1 LLM-based models

These models leverage LLMs to model student learning behavior by encoding question content and dialogue history to predict student responses.

- **LLMKT**: leverages textual question and conversation information in tutor–student dialogues using an LLM-based architecture, which is a previous state-of-the-art dialogue-based KT model (Scarlatos et al., 2025a).

### 4.2.2 Deep learning-based models

These models employ conventional deep neural networks to capture student knowledge states from sequences of their historical interactions, which consist of question tags, KC tags and binary responses, rather than LLMs.

- **DKT**: leverages an LSTM to sequentially model students’ learning processes and predict their future performance (Piech et al., 2015).
- **DKVMN**: leverages a static key memory matrix to encode relationships among KCs and a dynamic value memory matrix to predict students’ knowledge mastery levels (Zhang et al., 2017).
- **SAINT**: leverages a Transformer-based architecture to model students’ sequential learning interactions over time (Choi et al., 2020).
- **AKT**: leverages a monotonic attention mechanism to explicitly capture the temporal distance between current questions and students’ historical interactions (Ghosh et al., 2020).
- **simpleKT**: leverages a Rasch-inspired architecture to model question-specific variations and dot-product attention for modeling students’ sequential learning (Liu et al., 2023b).

## 4.3 Metrics and Experimental Setup

To quantitatively evaluate KT models in dialogue settings, we report Area Under the Curve (AUC) and Accuracy. We compute these metrics over all correctness labels except for the first label in each dialogue, since no prior observations are available for KT prediction at the initial step.

For our framework, we adopt Llama-3.1-8B-Instruct (Grattafiori et al., 2024) as the base model. We implement the model using the HuggingFace Transformers library (Wolf et al., 2020) and fine-tune it with LoRA (Hu et al.) on NVIDIA RTX L40 GPUs for efficient adaptation.

We perform a grid search to select hyperparameters for each model, choosing the checkpoint with the highest validation AUC to evaluate the test set. For LLM-based models, we search learning rates  $\in \{1 \times 10^{-5}, 5 \times 10^{-5}, 1 \times 10^{-4}, 2 \times 10^{-4}\}$  and LoRA ranks  $\in \{16, 32\}$  and fixing the LoRA scaling factor  $\alpha = 16$ . The effective batch size is 64 via gradient accumulation, with gradient clipping at 1.0. We train the models for 5 epochs using BFloat16 mixed precision. For other KT baselines, we search learning rates  $\in \{1 \times 10^{-4}, 2 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}, 2 \times 10^{-3}, 5 \times 10^{-3}\}$  and embedding sizes  $\in \{8, 16, 32, 64, 128, 256, 512\}$ . The batch size is 64 and the total epoch is 100, the remaining hyperparameters follow the default settings in pyKT (Liu et al., 2022). We train all models with AdamW optimizer with a weight decay of  $1 \times 10^{-2}$ .

## 5 Experimental Results

Model	QATD <sub>2k</sub>		MathDial	
	Accuracy	AUC	Accuracy	AUC
Deep Learning-based Models				
DKT	60.19	54.73	59.75	63.14
DKVMN	58.24	51.79	56.42	60.64
SAINT	60.71	53.76	59.80	63.22
AKT	60.88	54.39	59.75	63.37
simpleKT	61.62	56.24	58.94	61.90
LLM-based Models				
LLMKT	<u>64.24</u>	<u>64.89</u>	<u>68.64</u>	<u>75.99</u>
Ours	<b>64.29</b>	<b>65.25</b>	<b>68.82</b>	<b>76.59</b>

Table 1: Student dialogue turn correctness prediction performance on QATD<sub>2k</sub> and MathDial datasets. Best results are shown in **bold** and second-best results are underlined.

### 5.1 Quantitative Results

Table 1 shows the overall performance in terms of AUC and Accuracy for all models. We have the following observations: First, our proposed framework achieves the best overall performance and improves upon the previous state-of-the-art LLMKT model by explicitly modeling student ability, item difficulty, and combining them with IRT-based prediction. Specifically, our framework improves AUC from 64.89 to 65.25 (+0.36) on QATD<sub>2k</sub> and from 75.99 to 76.59 (+0.60) on MathDial. These improvements suggest that explicitly

modeling question difficulty provides additional useful information for monitoring student learning progress. Although the performance gains are relatively modest, our framework introduces interpretable outputs, which we further analyze model interpretability in Section 5.3; Second, in general, LLM-based methods consistently outperform deep learning-based KT models in terms of both AUC and Accuracy, demonstrating the advantage of leveraging pretrained LLMs in dialogue-based KT settings. We attribute this improvement to the ability of pretrained LLMs to model rich textual content in dialogues, enabling us to better track student learning after fine-tuning on response correctness labels. Third, among deep learning-based KT methods, DKT achieves comparable performance with those complex architectures such as SAINT and AKT in both datasets. These results suggest that increasing architectural complexity does not necessarily yield significant performance gains. Instead, simpler architectures may generalize more effectively in dialogue-based KT settings, particularly under limited-data conditions (fewer than 2,500 dialogues), where complex models are more susceptible to overfitting.

## 5.2 Qualitative Analysis: Difficulty Analysis

To show that the learned difficulty parameters are indeed meaningful and interpretable, we measure the consistency between predicted KC difficulty and “ground-truth” KC difficulty. We use QATD<sub>2k</sub> instead of MathDial for this analysis, since QATD<sub>2k</sub> is collected from real-world educational environments and therefore better captures authentic tutor-student behaviors. Similar to (Lee et al., 2024), we follow classical test theory to define the ground-truth KC difficulty based on empirical correctness statistics as  $1 - \frac{N_{\text{correct}}}{N_{\text{total}}}$ , i.e., the portion of incorrect responses on this KC (Brown, 2013). To ensure a fair comparison, we normalize the predicted difficulty values using min-max normalization. As illustrated in Figure 3, the fitted regression line follows a similar increasing trend to the perfect agreement line, indicating that the predicted difficulty generally increases as the ground-truth difficulty increases. To quantitatively support this observed trend, we compute the Pearson correlation between predicted and ground-truth difficulty values across all KCs (Benesty et al., 2009). Results show a moderate positive correlation ( $r = 0.368$ ,  $p < 0.001$ ), suggesting that the learned difficulty parameters capture meaningful




Figure 3: Comparison of predicted and ground-truth KC difficulty across KCs in the test set of the QATD<sub>2k</sub> dataset. Each point represents one KC. Point size is proportional to the number of samples associated with each KC. The dashed gray line indicates perfect agreement, while the solid blue line shows the fitted linear regression trend.

empirical difficulty patterns in real-world tutoring data across KCs.

## 5.3 “Learning Curve” Analysis

A common approach to evaluate KT quality is to examine whether predicted learning dynamics align with cognitive theories such as the power law of practice (Snoddy, 1926). Following prior work (Shi et al., 2023; Scarlatos et al., 2025a), we visualize learning trajectories of each KC. We use QATD<sub>2k</sub> instead of MathDial for this analysis, since QATD<sub>2k</sub> is collected from real-world educational environments and therefore better captures authentic tutor-student behaviors. Specifically, for each KC, we group student interactions by the number of prior opportunities to practice that KC, where an “occurrence” refers to each time the KC appears in a tutor-posed task within a dialogue. For each occurrence index, we compute the average predicted mastery across all corresponding dialogue turns in the test set. The resulting curves therefore illustrate how the estimated mastery evolves as students encounter the same KC multiple times.

We examine plots of the top three most frequent KCs, as they contain sufficient samples and provide the most intuitive visualization of learning trends. As shown in Figure 4, LLMKT, which directly treats predicted correctness probability as knowledge mastery, exhibits noticeable oscillations across KC occurrences. The predicted probability decreases even after repeated practice, suggesting unstable representations of student mastery. In con-

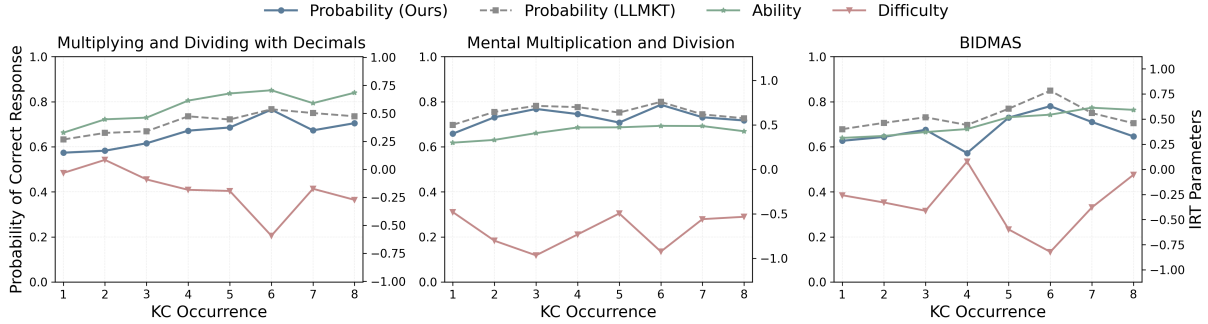


Figure 4: Learning trajectories for the three most frequent KCs in the QATD<sub>2k</sub> dataset.

trast, our framework separates ability estimation from correctness prediction, resulting in smoother and more stable learning trajectories. The estimated student ability consistently increases with repeated KC occurrences across all three KCs, reflecting an increase in knowledge mastery that aligns with the cognitive theories of gradual learning (Ritter and Schooler, 2002).

Furthermore, fluctuations in predicted correctness probability from our framework are largely associated with changes in tutor-posed task difficulty. For example, temporary drops in predicted correctness probability often coincide with increases in estimated task difficulty, while the underlying ability continues to increase. Such fluctuations are expected in dialogue-based tutoring settings, where task difficulty may decrease when the tutor provides hints or scaffolding, and increase when the tutor switches to a different, more difficult part of the question or presses the student for a complete solution. Hence, our framework reflects realistic interactions between student learning progress and task difficulty. Similar trends observed across multiple KCs further demonstrate that our framework produces more stable and cognitively plausible learning dynamics than probability-based mastery estimation models such as LLMKT. We also provide learning trend visualizations at the dialogue level in Appendix B.

#### 5.4 Qualitative Analysis: Case Study

We further conduct a qualitative case study to examine how the difficulty of tutor-posed tasks and the student’s knowledge state evolve in a dialogue in the QATD<sub>2k</sub> dataset. Specifically, we analyze one representative dialogue from the test set. As shown in Table 2, the tutor dynamically adjusts the difficulty of the posed tasks based on the student’s responses. When the student fails to answer

correctly, the tutor decomposes the original question into simpler sub-tasks, leading to a decrease in the estimated task difficulty, from 0.671 to 0.360. In contrast, when the student answers correctly, the tutor proceeds with follow-up questions that maintain or slightly increase the difficulty level, moving the dialogue forward. Meanwhile, the student’s knowledge state demonstrates a generally increasing trend across the dialogue. Notably, even when incorrect responses occur, the estimated student ability continues to improve gradually with the tutor’s guidance, from 0.224 to 0.628 across the full dialogue. This observation highlights that learning progress may still occur despite temporary errors, as scaffolding and intermediate feedback from the tutor help consolidate the student’s KC understanding. Overall, this case study illustrates that our framework captures realistic tutoring dynamics, where task difficulty adapts to student performance while student ability evolves steadily through guided interaction.

## 6 Related Work

### 6.1 Student Modeling in Dialogues

Recently, with the emergence of LLM-based tutoring systems, several studies have explored KT in dialogue settings. Scarlatos et al. (2025a) leveraged the textual content in dialogues to estimate student knowledge. Wang et al. (2025) proposed a workflow that integrates KT with turn-by-turn verification to improve coding tutoring (Wang et al., 2025). However, these approaches overlook the modeling of question difficulty, which plays a critical role in KT. On the other hand, our approach introduces a difficulty-aware dialogue-based KT framework that explicitly aligns the student knowledge state and tutor-posed task difficulty, enabling fine-grained modeling of student state evolution. Beyond KT, recent studies have increasingly ex-

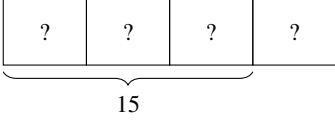
$q$	$t_j$	$d_j$	$s_j$	$\theta_j$	$r_j$
Which question could the bar model represent?  A: Calculate $\frac{3}{4}$ of 15 B: $\frac{3}{4}$ of $\square = 15$ C: Calculate $\frac{1}{3}$ of 15 D: $\frac{4}{3}$ of $\square = 15$	what is the bar model split up into?	–	15	0.224	0
	and in terms of the whole bar model, how much does the 15 cover?	0.671	half.	0.475	0
	not quite! How many boxes is the bar split into?	0.360	4	0.531	1
	okay good, and how many does the 15 span across?	0.447	All of the bar	0.605	0
	as in the bottom arrow. How many boxes is it under?	0.440	1	0.628	0
	not quite, can you see that the 15 and the arrow go across 3 boxes?	0.405	–	–	–

Table 2: Case study of an adaptive tutoring dialogue illustrating how task difficulty ( $d_j$ ) and student ability ( $\theta_j$ ) evolve over successive interactions in the QATD<sub>2k</sub> dataset. The tutor dynamically scaffolds the problem by decomposing it into simpler sub-tasks following incorrect responses, while progressively maintaining or increasing difficulty after correct answers, leading to a gradual improvement in the student’s knowledge state.

plored LLM-based student simulation in educational dialogue settings. Prior work has demonstrated that LLMs can be used to simulate realistic student learning behaviors, including the development of learning curves and misconception patterns (Schmucker et al., 2024; Jin et al., 2024). Other studies simulate student behaviors conditioned on demographic, behavioral, and personality traits (Markel et al., 2023; Li et al., 2025a), as well as varying learning goals and trajectories (Sharma and Li, 2024).

## 6.2 Traditional Knowledge Tracing

Traditional KT models capture a student’s learning behavior in a sequence of her/his historical interactions, which consist of question tags, KC tags and binary responses (Chen et al., 2023; Liu et al., 2023a; Huang et al., 2023, 2026; Liu et al., 2025; Guo et al., 2026). For example, Piech et al. (2015) introduced the first DKT model, utilizing an LSTM to estimate knowledge mastery. Zhang et al. (2017) proposed a static key memory matrix to store KC relationships and a dynamic value memory matrix for predicting knowledge mastery levels. To incorporate the difficulty factor to enhance the knowledge state estimation, Ghosh et al. (2020) introduced a scalar difficulty parameter to control how far a question deviates from the underlying concept during question representation learning. Shen et al. (2022) established the relationship between student knowledge state and question difficulty level to improve KT performance. Liu et al. (2024) replaced question tags with difficulty levels to facilitate the translation of model predictions into interpretable

skill-level knowledge states, further highlighting the importance of difficulty modeling in KT.

## 7 Conclusion and Future Work

In this work, we introduce an interpretable difficulty-aware conversational KT framework that jointly models student ability and question difficulty in tutor-student dialogues. Our framework aims to capture fine-grained difficulty signals and provide more interpretable predictions of student performance, connecting LLM-based modeling with cognitively motivated assumptions. Quantitative results show modest improvements over existing KT baselines, suggesting the potential benefit of explicitly modeling both student ability and task difficulty. Empirical qualitative analyses provide further evidence for the interpretability of the learned parameters and indicate the value of incorporating an IRT predictor. In particular, the learned difficulty parameters show positive correlations with empirical difficulty statistics, suggesting that the model captures meaningful and cognitively plausible difficulty patterns across KCs. Additionally, the estimated ability trajectories often exhibit smooth growth patterns broadly consistent with gradual learning, while fluctuations in predicted correctness tend to correspond to changes in tutor-proposed task difficulty.

There are many avenues for future work. First, since we explicitly consider the difficulty factor in dialogue-based KT settings, we will further explore more effective methods for capturing fine-grained difficulty signals from tutoring dialogues and investigate how task difficulty evolves dynamically

during tutoring interactions. Second, for student knowledge state estimation, we plan to leverage advanced reasoning capabilities and reinforcement learning techniques to improve the accuracy and robustness of the estimation. In particular, future work will explore how multi-step reasoning can ensure consistency across dialogue turns and how reinforcement learning methods can refine estimation strategies based on long-term prediction performance. Third, we plan to explore misconception-aware conversational KT by detecting and tracking students' misconceptions from dialogue interactions. By integrating misconception signals with estimated knowledge states and difficulty parameters, the framework may provide more fine-grained and interpretable student performance prediction.

### Acknowledgments

This work is partially supported by Renaissance Philanthropy under the learning engineering virtual institute (LEVI) initiative and the NSF under grant 2237676.

### Limitations

There are several practical limitations to our work. First, since our work is the first to explicitly estimate the difficulty of tutor turns in dialogues, there are no existing difficulty prediction baselines we can compare against. Second, we only conduct experiments on math dialogues, and do not investigate other domains such as language learning or computer science. However, this is largely because KT in dialogues is currently only well-established in math. Finally, while our framework and other LLM-based approaches are significantly more accurate than traditional KT approaches, they are much more computationally expensive too. An important line of future work will be improving the efficiency of LLM-based approaches while maintaining high predictive accuracy.

### Ethical Considerations

There are both potential societal benefits and risks associated with our work. Accurately modeling student ability and difficulty in dialogues is a necessary step to facilitating automated assessments in the dialogue setting. This can potentially benefit students through more accurate assessment and subsequently more targeted learning actions. However, as is the case with many other advances in educational AI, further automating assessments creates

a risk of replacing teachers' jobs. Our intention is for educational AI tools such as ours to be used in conjunction with teachers to improve outcomes for students, rather than be used to replace teachers. Finally, there is a risk of bias in our framework, as is common in all AI methods, due to inherent biases in training data. For instance, it is possible that students from populations that are less represented in training data are given less accurate assessments, therefore risking worse outcomes for those students. It is important that educational AI tools are thoroughly screened for bias before deployment with real students.

### References

- John R Anderson. 2000. *Learning and memory: An integrated approach*. John Wiley & Sons.
- Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. In *Noise reduction in speech processing*, pages 1–4. Springer.
- Allan Birnbaum. 1968. Some latent trait models and their use in inferring an examinee's ability. *Statistical theories of mental test scores*.
- James Dean Brown. 2013. Classical test theory. In *The Routledge handbook of language testing*, pages 337–349. Routledge.
- Carnegie Learning. 2024. Livehint overview. <https://support.carnegielearning.com/help-center/math/livehint/article/livehint-overview/>. Accessed: 2026-03-18.
- Jiahao Chen, Zitao Liu, Shuyan Huang, Qiongqiong Liu, and Weiqi Luo. 2023. Improving interpretability of deep sequential knowledge tracing models with question-centric cognitive representations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 14196–14204.
- Youngduck Choi, Youngnam Lee, Junghyun Cho, Jineon Baek, Byungsoo Kim, Yeongmin Cha, Dongmin Shin, Chan Bae, and Jaewe Heo. 2020. Towards an appropriate query, key, and value computation for knowledge tracing. In *Proceedings of the seventh ACM conference on learning@ scale*, pages 341–344.
- Aritra Ghosh, Neil Heffernan, and Andrew S Lan. 2020. Context-aware attentive knowledge tracing. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2330–2339.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

- Teng Guo, Bojun Zhan, Shuyan Huang, Jiahao Chen, Xiangyu Zhao, Mingliang Hou, and Zitao Liu. 2026. Improving knowledge tracing through multi-source scaling with decoder-only transformers. *ACM Transactions on Intelligent Systems and Technology*, 17(3):1–25.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Shuyan Huang, Zitao Liu, Qiongqiong Liu, Jiahao Chen, and Yaying Huang. 2026. Dual-attentional time-aware fusion networks for knowledge tracing. *Information Fusion*, 126:103573.
- Shuyan Huang, Zitao Liu, Xiangyu Zhao, Weiqi Luo, and Jian Weng. 2023. Towards robust knowledge tracing models via k-sparse attention. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*, pages 2441–2445.
- Hyounghook Jin, Seonghee Lee, Hyungyu Shin, and Juho Kim. 2024. Teach ai how to code: Using large language models as teachable agents for programming education. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–28.
- Khan Academy. 2023. Supercharge your teaching experience with khanmigo. <https://www.khanmigo.ai/>. Accessed: 2026-03-18.
- Jaewook Lee, Alexander Scarlatos, Simon Woodhead, and Andrew Lan. 2026. Letting tutor personas speak up for llms: Learning steering vectors from dialogue via preference optimization. *arXiv preprint arXiv:2602.07639*.
- Unggi Lee, Sungjun Yoon, Joon Seo Yun, Kyoungsoo Park, YoungHoon Jung, Damji Stratton, and Hyeoncheol Kim. 2024. Difficulty-focused contrastive learning for knowledge tracing with a large language model-based difficulty prediction. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4891–4900.
- Haoxuan Li, Jifan Yu, Xin Cong, Yang Dang, Daniel Zhang-Li, Lu Mi, Yisi Zhan, Huiqin Liu, and Zhiyuan Liu. 2025a. Which type of students can llms act? investigating authentic simulation with graph-based human-ai collaborative system. *arXiv preprint arXiv:2502.11678*.
- Haoxuan Li, Jifan Yu, Yuanxin Ouyang, Zhuang Liu, Wenge Rong, Huiqin Liu, Juanzi Li, and Zhang Xiong. 2025b. Explainable few-shot knowledge tracing. *Frontiers of Digital Education*, 2(4):34.
- Guimei Liu, Huijing Zhan, and Jung-jae Kim. 2024. Question difficulty consistent knowledge tracing. In *Proceedings of the ACM web conference 2024*, pages 4239–4248.
- Zitao Liu, Shuyan Huang, Teng Guo, Mingliang Hou, and Qianru Liang. 2025. A prompt-driven framework for multi-domain knowledge tracing. *Machine Learning*, 114(4):87.
- Zitao Liu, Qiongqiong Liu, Jiahao Chen, Shuyan Huang, Boyu Gao, Weiqi Luo, and Jian Weng. 2023a. Enhancing deep knowledge tracing with auxiliary tasks. In *Proceedings of the ACM web conference 2023*, pages 4178–4187.
- Zitao Liu, Qiongqiong Liu, Jiahao Chen, Shuyan Huang, and Weiqi Luo. 2023b. simplekt: a simple but tough-to-beat baseline for knowledge tracing. *The Eleventh International Conference on Learning Representations*.
- Zitao Liu, Qiongqiong Liu, Jiahao Chen, Shuyan Huang, Jiliang Tang, and Weiqi Luo. 2022. pykt: a python library to benchmark deep learning based knowledge tracing models. *Advances in Neural Information Processing Systems*, 35:18542–18555.
- Jakub Macina, Nico Daheim, Ido Hakimi, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2025. Math-tutorbench: A benchmark for measuring open-ended pedagogical capabilities of llm tutors. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 204–221.
- Julia M Markel, Steven G Opferman, James A Landay, and Chris Piech. 2023. Gpteach: Interactive training with gpt-based students. In *Proceedings of the tenth acm conference on learning@ scale*, pages 226–236.
- Allen Newell and Paul S Rosenbloom. 1981. Mechanisms of skill acquisition and the law of practice. *Cognitive skills and their acquisition*, pages 1–55.
- Andre Nickow, Philip Oreopoulos, and Vincent Quan. 2020. The impressive effects of tutoring on prek-12 learning: A systematic review and meta-analysis of the experimental evidence.
- Benjamin D Nye, Arthur C Graesser, and Xiangen Hu. 2014. Autotutor and family: A review of 17 years of natural language tutoring. *International Journal of Artificial Intelligence in Education*, 24(4):427–469.
- Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. 2015. Deep knowledge tracing. *Advances in neural information processing systems*, 28.
- Georg Rasch. 1993. *Probabilistic models for some intelligence and attainment tests*. ERIC.
- Frank E. Ritter and Lael J. Schooler. 2002. The learning curve. In *International Encyclopedia of the Social and Behavioral Sciences*. Pergamon.
- Alexander Scarlatos, Ryan S Baker, and Andrew Lan. 2025a. Exploring knowledge tracing in tutor-student dialogues using llms. In *Proceedings of the 15th International Learning Analytics and Knowledge Conference*, pages 249–259.

- Alexander Scarlatos, Nigel Fernandez, Christopher Ormerod, Susan Lottridge, and Andrew Lan. 2025b. Smart: Simulated students aligned with item response theory for question difficulty prediction. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25082–25105.
- Alexander Scarlatos, Jaewook Lee, Simon Woodhead, and Andrew Lan. 2026. [Simulated students in tutoring dialogues: Substance or illusion?](#) Preprint, arXiv:2601.04025.
- Alexander Scarlatos, Naiming Liu, Jaewook Lee, Richard Baraniuk, and Andrew Lan. 2025c. Training llm-based tutors to improve student learning outcomes in dialogues. In *International Conference on Artificial Intelligence in Education*, pages 251–266. Springer.
- Timo Schick and Hinrich Schütze. 2021. Exploiting cloze-questions for few-shot text classification and natural language inference. In *Proceedings of the 16th conference of the European chapter of the association for computational linguistics: main volume*, pages 255–269.
- Robin Schmucker, Meng Xia, Amos Azaria, and Tom Mitchell. 2024. Ruffle&riley: Insights from designing and evaluating a large language model-based conversational tutoring system. In *International Conference on Artificial Intelligence in Education*, pages 75–90. Springer.
- Paras Sharma and Qichang Li. 2024. Designing simulated students to emulate learner activity data in an open-ended learning environment. In *Proceedings of the 17th International Conference on Educational Data Mining*, pages 986–989.
- Shuanghong Shen, Zhenya Huang, Qi Liu, Yu Su, Shijin Wang, and Enhong Chen. 2022. Assessing student’s dynamic knowledge state by exploring the question difficulty effect. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, pages 427–437.
- Yang Shi, Robin Schmucker, Min Chi, Tiffany Barnes, and Thomas Price. 2023. Kc-finder: Automated knowledge component discovery for programming problems.
- George S Snoddy. 1926. Learning and stability: a psychophysiological analysis of a case of motor learning with clinical applications. *Journal of Applied Psychology*, 10(1):1.
- Jian Wang, Yinpei Dai, Yichi Zhang, Ziqiao Ma, Wenjie Li, and Joyce Chai. 2025. Training turn-by-turn verifiers for dialogue tutoring agents: The curious case of llms as your coding tutors. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 12416–12436.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and 1 others. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.
- Matthew Zent, Digory Smith, and Simon Woodhead. 2025. Piivot: A lightweight nlp anonymization framework for question-anchored tutoring dialogues. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 27467–27476.
- Jiani Zhang, Xingjian Shi, Irwin King, and Dit-Yan Yeung. 2017. Dynamic key-value memory networks for knowledge tracing. In *Proceedings of the 26th international conference on World Wide Web*, pages 765–774.

## A Prompts

### A.1 Prompt for Knowledge Estimation

*You are an experienced math teacher. Given a dialogue where the student is working through a multiple-choice diagnostic math question, based on the student’s responses and demonstrated understanding in the dialogue, classify the student’s current ability level. Respond with exactly one token from: GOOD BAD.*

[BEGIN QUESTION]

Question Text

[END QUESTION]

[BEGIN DIALOGUE]

Dialogue Context

[END DIALOGUE]

Specifically, the input prompt  $I^K$  consists of the following three key components:

- **Task Definition:** A natural language instruction that defines the task and specifies the expected output format for the LLM.
- **Question Text:** The diagnostic question presented to the student, typically enclosed within structured markers [BEGIN QUESTION] and [END QUESTION].
- **Dialogue Context:** The historical teacher–student interaction associated with the question, enclosed within markers [BEGIN DIALOGUE] and [END DIALOGUE].




Figure 5: Learning trajectories for dialogues in the QATD<sub>2k</sub> dataset.

## A.2 Prompt for Difficulty Estimation

*You are an experienced math teacher. Given a dialogue where the student is working through a multiple-choice diagnostic math question, based on the question content and required knowledge components, classify the question’s difficulty level. Respond with exactly one token from: HARD EASY.*

[BEGIN QUESTION]  
 Question Text  
 [END QUESTION]

[BEGIN DIALOGUE]  
 Dialogue Context  
 [END DIALOGUE]

[BEGIN CURRENT TEACHER TURN]  
 Teacher-posed Task  
 [END CURRENT TEACHER TURN]

[BEGIN KC]  
 KC Names  
 [END KC]

Specifically, the input prompt  $I^D$  consists of the following two key components:

- **Teacher-posed Task:** The next teacher-posed question that the student is expected to respond to is enclosed within markers [BEGIN CURRENT TEACHER TURN] and [END CURRENT TEACHER TURN], representing the immediate task used for predicting the student’s knowledge state.
- **KC Names:** The associated KCs with the next tutor-posed task enclosed within markers [BEGIN KC] and [END KC]. The model must infer

from prior interactions whether the student has mastered the required knowledge.

## B Additional “Learning Curve” Analysis

We provide learning curve analysis based on dialogue level in Figure 5.