

# Exploring LLMs for Predicting Tutor Strategy and Student Outcomes in Dialogues

Fareya Ikram, Alexander Scarlatos, Andrew Lan

University of Massachusetts Amherst  
{fikram, ajscarlatos, andrewlan}@umass.edu

## Abstract

Tutoring dialogues have gained significant attention in recent years, given the prominence of online learning and the emerging tutoring abilities of artificial intelligence (AI) agents powered by large language models (LLMs). Recent studies have shown that the strategies used by tutors can have significant effects on student outcomes, necessitating methods to predict how tutors will behave and how their actions impact students. However, few works have studied predicting tutor strategy in dialogues. Therefore, in this work we investigate the ability of modern LLMs, particularly Llama 3 and GPT-4o, to predict both future tutor moves and student outcomes in dialogues, using two math tutoring dialogue datasets. We find that even state-of-the-art LLMs struggle to predict future tutor strategy while tutor strategy is highly indicative of student outcomes, outlining a need for more powerful methods to approach this task.

## 1 Introduction

Tutoring has been shown to be highly effective in increasing student learning, both when administered by human tutors or intelligent tutoring systems (Nickow et al., 2020; Nye et al., 2014). Recently, several automated tutors, powered by large language models (LLMs), have been deployed in educational settings, such as Khan Academy’s Khanmigo (Khan Academy, 2023) or Carnegie Learning’s LiveHint (Carnegie Learning, 2024). To ensure that students benefit from these tools, it is important to study the strategies used by tutors and how they impact student learning outcomes.

Tutor strategy is commonly formalized using “moves”, or high-level pedagogical actions taken in any given dialogue turn to support student learning (Demszky and Hill, 2023; Macina et al., 2023; Suresh et al., 2022). Recent studies have shown that explicit move and strategy information can be used to improve tutor effectiveness (Wang et al.,

2024a,b). Others that train LLMs to be effective tutors (Tack et al., 2023; Team et al., 2024; Sonkar et al., 2024; Huber et al., 2023; Vasselli et al., 2023; Scarlatos et al., 2025b) have also highlighted the importance of pedagogical strategy. Several prior works have used tutor moves to predict student outcomes (Lin et al., 2022; Borchers et al., 2024; Abdelshiheed et al., 2024; Yin et al., 2025), though text alone processed by LLMs is often sufficient (Scarlatos et al., 2025a; Chen et al., 2024). In this work, we explicitly study the effect of move annotations, compared to text alone, on predicting tutor strategy and student outcomes.

While many works have studied how to *identify* tutor moves in dialogues (Demszky et al., 2021; Wang and Demszky, 2024; Moreau-Pernet et al., 2024; Wang et al., 2023; McNichols and Lan, 2025), there are few works studying how to predict *future* tutor moves. One prior work does so using GRUs and RoBERTa (Ganesh et al., 2021), though to the best of our knowledge, none have studied how generative language models, such as Meta’s Llama 3 (Dubey et al., 2024) or OpenAI’s GPT-4 (OpenAI, 2024b), perform on this task.

In this work, to address the needs of understanding tutor strategy and its effect on student outcomes, we seek to answer the following research questions: **RQ1**: Can LLMs predict *tutor strategy* using tutor moves and dialogue history?, **RQ2**: Can LLMs predict *student outcomes* using tutor moves and dialogue history?, and **RQ3**: Which tutoring strategies have the highest impact on student outcomes? To the best of our knowledge, our study is the first to jointly address these questions using modern generative LLMs. Overall, we find that tutor strategy prediction is highly challenging, with student outcome prediction being easier and facilitated by tutor move annotations. Our findings indicate that tutor strategy prediction is an important and challenging task worth further study in future work.

Speaker	Utterance
Tutor	Hi Ayisha, please talk me through your solution ( <b>generic</b> )
Student	I started by noting .... I concluded that the desk cost her \$350, since that was her winning bid. ....
Student	Yes, you're right...Carmen's bid added an additional \$150, making the total \$200 + \$150 + \$150 = \$350.
Tutor	Check your calculation of \$200 + \$150 + \$150 = \$350. Your total is not correct ( <b>probing</b> )

Table 1: A tutor-student dialogue from MathDial, showing annotated moves for tutor turns.

## 2 Methodology

In this section, we detail the four primary tasks we use to study tutor strategy and student outcome prediction, as well as the various LLM and non-LLM methods we evaluate on these tasks. First, we define our notation. Each tutor-student dialogue,  $D = \{T_i, S_i, M_i, E_i, C\}_{i=1}^N$ , contains a sequence of alternating textual student utterances  $S_i$  and tutor utterances  $T_i$ . Each tutor turn is associated with one or more categorical moves, i.e., granular pedagogical actions,  $M_i$ , and each student turn is optionally associated with a binary measure of success  $E_i$ . Each dialogue is also labeled with a final binary measure of success,  $C$ . As our focus is to study tutor strategy and student outcomes, rather than student behavior, we do not use student move labels. We show an example dialogue in Table 1.

### 2.1 Tasks

We now detail our four primary tasks. First, we examine future tutor move prediction to investigate if models can predict tutoring strategy. Second, we examine tutor move classification to investigate if models can infer moves from tutor utterances. Third, we examine future dialogue success prediction to investigate if models can infer the outcome of a dialogue from limited context. Finally, we examine next student turn success prediction to investigate if models can predict short-term student outcomes. We formalize these tasks as follows:

$$\hat{M}_{i+1} = f_{\theta}(\{T_j, S_j, M_j\}_{j=1}^i) \quad (1)$$

$$\hat{M}_i = f_{\theta}(\{T_j, S_j, M_j\}_{j=1}^{i-1}, T_i, S_i) \quad (2)$$

$$\hat{C}_i = f_{\theta}(\{T_j, S_j, M_j\}_{j=1}^i) \quad (3)$$

$$\hat{E}_{i+1} = f_{\theta}(\{T_j, S_j, M_j\}_{j=1}^i) \quad (4)$$

## 2.2 LLM-Based Methods

We evaluate tutor move and student outcome prediction using two large language model (LLM)-based methods. First, we fine-tune the **Llama 3.2 3B** model using Low-Rank Adaptation (LoRA) (Saari et al., 2018), where we train the model to predict the labels as text tokens following the input, using comma separation for multi-label turns. Second, we use zero-shot prompting with **GPT-4o**, where we prompt the model to follow the annotation schema in each particular dataset to identify and predict tutor moves. We do not use GPT-4o for student outcome prediction because pre-trained LLMs do not show high alignment with student behavior (Liu et al., 2025). We show an example prompt in Figure 7, and provide further implementation details in Appendix A. In order to determine how tutor move information impacts LLM predictions, for both methods, we experiment with including only the dialogue as input, as well as both the dialogue and previous turn move labels as input.

### 2.3 Baselines

We additionally experiment with three traditional baselines, where the input space only uses the move labels of previous turns. First, we employ a second-order **Markov Chain** (Boyer et al., 2009), which estimates the probability of a tutor move or student success given the two preceding moves. Next, we employ **Logistic Regression**, using the frequency distribution of moves up to the current turn as input features. Finally, in order to capture temporal dependencies beyond adjacent moves, we employ an **LSTM** model on sequences of tutor move types encoded as multi-hot vectors.

We do not use these baselines for current move identification since they do not process text. Additionally, we do not use Markov Chain and Logistic Regression for predicting future moves for AlgebraNation because it is a multi-label task; the input space for Markov Chain would be exponentially large, and Logistic Regression would suffer from an overwhelming amount negative labels per class.

## 3 Experiments

### 3.1 Datasets

We experiment with two math tutoring datasets, MathDial (Macina et al., 2023) and AlgebraNation (Lyu et al., 2024), to answer our research questions. MathDial contains one-on-one dialogues where a tutor guides a student through solving multi-step

Model	MathDial				AlgebraNation			
	Tutor Move		Future Tutor Move		Tutor Move		Future Tutor Move	
	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1
Markov Chain	–	–	42.27	<u>42.81</u>	–	–	–	–
Logistic Regression	–	–	44.69	34.58	–	–	–	–
LSTM	–	–	<u>46.19</u>	31.13	–	–	3.36	23.17
GPT-4o - Dialogue	49.56	48.76	36.48	32.49	<u>79.22</u>	54.30	<u>65.72</u>	<b>27.32</b>
GPT-4o - Dialogue & Moves	50.07	<u>49.03</u>	31.72	29.00	<b>83.82</b>	57.40	<b>69.86</b>	<u>26.33</u>
Llama 3 - Dialogue	<u>52.58</u>	45.88	42.74	32.92	58.82	<u>62.78</u>	24.53	20.20
Llama 3 - Dialogue & Moves	<b>59.69</b>	<b>57.26</b>	<b>50.35</b>	<b>49.33</b>	63.42	<b>68.90</b>	25.69	21.09

Table 2: Results for identifying tutor moves and predicting future tutor moves. Llama 3 performs best on MathDial, while GPT-4o generally performs best on AlgebraNation.

math reasoning problems. The student utterances in this dataset are simulated by an LLM, prompted to mimic common student misconceptions. Each tutor turn is labeled with one of four moves: *probing*, *focus*, *telling*, and *generic*. We leverage turn-level student correctness labels from (Scarlatos et al., 2025a). Our final dataset contains 2,484 dialogues with a 1,947/537 train/test split.

AlgebraNation contains discourse from an online forum where students pose questions and discuss with both tutors and peers. Each tutor turn is labeled with any number of 16 move types. Each post on the forum is marked with success if the responses resolve the original student question. The dataset contains 2,318 forum posts, which we split into a 1,854/464 train/test split. We show label distributions for both datasets in Appendix C.

### 3.2 Evaluation Metrics

We employ two widely used metrics: i) accuracy (**Acc.**), the portion of predicted labels that match the ground truth, and ii) weighted **F1**, the harmonic mean of precision and recall, weighted by label frequency to account for imbalanced label distributions. For AlgebraNation, where each tutor turn may have multiple moves, we use exact match across all moves in a turn to compute accuracy.

### 3.3 Tutor Move Prediction

**Quantitative Analysis** We show the results for tutor move prediction Table 2. Across all methods, predicting the next tutor move proves to be more difficult than classifying the current move, particularly for AlgebraNation, which contains real-world interactions and more granular move definitions. LLMs improve over baselines for future move prediction, showing the importance of textual context and powerful models for this task. However, the F1 for future move prediction is low overall, only

reaching 49% for MathDial and 27% for AlgebraNation. These results indicate that tutor behavior is highly unpredictable, and that even state-of-the-art LLMs struggle to predict future tutor moves.

Additionally, the results across models and datasets are inconsistent; Llama 3 performs better on MathDial while GPT-4o performs better on AlgebraNation. Notably, unlike the other methods, GPT-4o was not trained on AlgebraNation, which exhibits a highly skewed label distribution (Figure 4). This imbalance may help explain why Llama 3 tends to default to predicting the majority class in the future move prediction task (Figure 2), a pattern not observed with GPT-4o. On the other hand, GPT-4o’s move prediction on MathDial likely suffers from confusion between label definitions, as we discuss in the qualitative analysis.

For the move identification task, Llama 3 performs best on MathDial, with the inclusion of annotated move labels significantly increasing performance. Similar to future move prediction, we attribute Llama 3’s higher performance to GPT-4o’s confusion between labels. For AlgebraNation, GPT-4o outperforms Llama 3 in accuracy but underperforms it in F1. This disparity can be explained by observing that Llama 3’s output distribution more closely resembles the ground truth distribution, as seen in Figure 1.

**Qualitative Analysis** We examine label misclassifications to investigate error patterns in model predictions, revealing dataset-specific challenges. For move classification, in MathDial, confusion frequently arises between *probing* and *focus* moves, while in AlgebraNation, *giving instruction* is often mistaken for *giving explanation* (Table 4). These misclassifications are likely attributable to conceptual overlap in the definitions of these categories, underscoring the nuanced nature of interpreting tu-

Model	MathDial				AlgebraNation	
	Turn Success		Dialogue Success		Dialogue Success	
	Acc.	F1	Acc.	F1	Acc.	F1
Markov Chain	53.75	<b>52.86</b>	53.27	49.01	62.99	63.55
Logistic Regression	52.80	<u>52.54</u>	64.05	66.43	69.30	63.24
LSTM	52.20	<u>52.27</u>	79.13	69.90	<u>77.98</u>	<u>77.76</u>
Llama - Dialogue	<u>63.56</u>	50.54	<b>79.21</b>	<b>70.11</b>	75.64	75.52
Llama - Dialogue & Moves	<b>64.11</b>	49.27	<u>79.16</u>	<u>69.96</u>	<b>81.00</b>	<b>80.84</b>

Table 3: Results for predicting student outcomes from dialogues. Baselines are competitive on MathDial, while Llama 3 performs best on AlgebraNation.

tor intentions. The MathDial authors also note that annotators had difficulty differentiating between *probing* and *focus* moves (Macina et al., 2023). On the other hand, the AlgebraNation annotation guidelines are more specific and instruction-driven, likely helping GPT-4o’s performance due to its ability to generalize well under clear, directive annotation schemes (OpenAI, 2024a). Notably, misclassifications decrease when observing previous move labels, reflected in Table 2, with these labels likely acting as informative in-context examples.

### 3.4 Student Outcome Prediction

**Quantitative Analysis** We show the results for student outcome prediction in Table 3. Notably, both Llama 3 and LSTM are able to achieve high performance on dialogue success prediction for both datasets, indicating the tractability of predicting near-term student outcomes in dialogues. However, we note that the MathDial results are inflated as they reflect majority class prediction, as shown in Figure 3. On the other hand, the distribution is more balanced for AlgebraNation, indicating that the outcomes of real students are more reliably predicted than the outcomes of simulated ones. We also see that previous move labels improve performance for Llama 3, showing that tutor moves complement dialogue text to infer student outcomes.

Predicting student outcomes at the turn-level proves to be more difficult than at the dialogue-level in MathDial, with baselines performing close to random chance. However, using LLMs improves performance on this task, capturing nuanced details in the dialogue text to help predict student behavior, as noted in (Scarlatos et al., 2025a).

**Regression Analysis** To investigate the impact of tutor moves on student outcomes, we examine the learned coefficients of our logistic regression model when predicting dialogue-level success and perform a Chi-squared analysis, shown in Tables

8 and 9. For AlgebraNation, *confirmatory feedback*, *giving instruction*, and *giving explanation* have the greatest positive impact on success. These tutor moves share a common thread: they are all instructionally supportive behaviors that actively guide the student’s understanding or progress. Each move either reinforces correct reasoning (confirmatory feedback), clarifies procedural steps (giving instruction), or deepens conceptual understanding (giving explanation). This correlation suggests that successful dialogues are those in which tutors take an active and supportive role in scaffolding the student’s learning process. For MathDial, *generic* and *probing* have the strongest positive impact on success, whereas *telling* has a negative impact on success. This finding aligns with prior work (Berghmans et al., 2014) showing that facilitative peer tutoring is more effective than directive tutoring.

## 4 Conclusion

In this work, we investigate the abilities and limitations of LLMs in classifying and predicting tutoring strategies and student outcomes in dialogues. We find that while results vary across models and datasets, LLMs outperform traditional baselines while still struggling at the task of tutor strategy prediction overall. Additionally, we find that student outcome prediction is tractable for LLMs, with tutor move information improving accuracy. Our findings emphasize the importance and challenges of studying tutor strategy in dialogues, given the impact that such strategies can have on student outcomes. Future work should explore how to improve tutor strategy prediction, potentially using in-context learning (Lee et al., 2024) or reinforcement learning (Li et al., 2024). Additionally, future work should explore how to suggest optimal tutor moves, potentially using reinforcement learning guided by student outcomes (Scarlatos et al., 2025b).



## Limitations

We identify several technical and practical limitations of our work. First, the generalizability of our results is constrained by the scope and nature of the datasets. MathDial involves synthetic student responses generated by LLMs, which may not reflect the complexity and variability of authentic student behavior. Conversely, AlgebraNation, while comprising real-world interactions, has a highly imbalanced label distribution that poses challenges for model evaluation. Additionally, our evaluation methodology predominantly relies on exact match accuracy and weighted F1 scores. These standard metrics may not fully capture the nuanced characteristics of our models. Finally, the absence of student move tracking in our current modeling approach may affect the results, as sequential modeling of student behavior could potentially enhance predictive performance.

## Acknowledgements

This work is partially supported by Renaissance Philanthropy via the learning engineering virtual institute (LEVI) and NSF grants 2118706, 2237676, and 2341948.

## References

- Mark Abdelshiheed, Jennifer K. Jacobs, and Sidney K. D’Mello. 2024. Aligning tutor discourse supporting rigorous thinking with tutee content mastery for predicting math achievement. In *Artificial Intelligence in Education*, pages 150–164, Cham. Springer Nature Switzerland.
- Inneke Berghmans, Lotte Michiels, Sara Salmon, Filip Dochy, and Katrien Struyven. 2014. Directive versus facilitative peer tutoring? a view on students’ appraisal, reported learning gains and experiences within two differently-tutored learning environments. *Learning Environments Research*, 17:437–459.
- Conrad Borchers, Kexin Yang, Jionghao Lin, Nikol Rummel, Kenneth R. Koedinger, and Vincent Aleven. 2024. [Combining dialog acts and skill modeling: What chat interactions enhance learning rates during ai-supported peer tutoring?](#) In *Proceedings of the 17th International Conference on Educational Data Mining*, pages 117–130, Atlanta, Georgia, USA. International Educational Data Mining Society.
- Kristy Elizabeth Boyer, Eun Young Ha, Michael D Wallis, Robert Phillips, Mladen A Vouk, and James C Lester. 2009. Discovering tutorial dialogue strategies with hidden markov models. In *Artificial Intelligence in Education*, pages 141–148. IOS Press.
- Carnegie Learning. 2024. Livehint overview. Online: <https://support.carnegielearning.com/help-center/math/livehint/article/livehint-overview/>.
- Jiahao Chen, Zitao Liu, Mingliang Hou, Xiangyu Zhao, and Weiqi Luo. 2024. [Multi-turn classroom dialogue dataset: Assessing student performance from one-on-one conversations](#). In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM ’24*, page 5333–5337, New York, NY, USA. Association for Computing Machinery.
- Dorottya Demszky and Heather Hill. 2023. The ncte transcripts: A dataset of elementary math classroom transcripts. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 528–538.
- Dorottya Demszky, Jing Liu, Zid Mancenido, Julie Cohen, Heather Hill, Dan Jurafsky, and Tatsunori B Hashimoto. 2021. Measuring conversational uptake: A case study on student-teacher interactions. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1638–1653.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Ananya Ganesh, Martha Palmer, and Katharina Kann. 2021. [What would a teacher do? Predicting future talk moves](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4739–4751, Online. Association for Computational Linguistics.
- Thomas Huber, Christina Niklaus, and Siegfried Handschuh. 2023. [Enhancing educational dialogues: A reinforcement learning approach for generating AI teacher responses](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 736–744, Toronto, Canada. Association for Computational Linguistics.
- Khan Academy. 2023. Supercharge your teaching experience with khanmigo. Online: <https://www.khanmigo.ai/>.
- Seanie Lee, Jianpeng Cheng, Joris Driesen, Alexandru Coca, and Anders Johannsen. 2024. [Effective and efficient conversation retrieval for dialogue state tracking with implicit text summaries](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 96–111, Mexico City, Mexico. Association for Computational Linguistics.

- Hang Li, Tianlong Xu, Jiliang Tang, and Qingsong Wen. 2024. [Knowledge tagging system on math questions via llms with flexible demonstration retriever](#). *Preprint*, arXiv:2406.13885.
- Jionghao Lin, Shaveen Singh, Lele Sha, Wei Tan, David Lang, Dragan Gašević, and Guanliang Chen. 2022. Is it a good move? mining effective tutoring strategies from human–human tutorial dialogues. *Future Generation Computer Systems*, 127:194–207.
- Naiming Liu, Shashank Sonkar, and Richard G. Baraniuk. 2025. [Do llms make mistakes like students? exploring natural alignment between language models and human error patterns](#). *Preprint*, arXiv:2502.15140.
- Bailing Lyu, Chenglu Li, Hai Li, Wangda Zhu, and Wanli Xing. 2024. Explaining technical, social, and discursive participation in online mathematical discussions. *Distance Education*, pages 1–24.
- Jakub Macina, Nico Daheim, Sankalan Chowdhury, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023. MathDial: A dialogue tutoring dataset with rich pedagogical properties grounded in math reasoning problems. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5602–5621, Singapore. Association for Computational Linguistics.
- Hunter McNichols and Andrew Lan. 2025. [The studychat dataset: Student dialogues with chatgpt in an artificial intelligence course](#). *Preprint*, arXiv:2503.07928.
- Baptiste Moreau-Pernet, Yu Tian, Sandra Sawaya, Peter Foltz, Jie Cao, Brent Milne, and Thomas Christie. 2024. [Classifying tutor discursive moves at scale in mathematics classrooms with large language models](#). In *Proceedings of the Eleventh ACM Conference on Learning @ Scale, L@S '24*, page 361–365, New York, NY, USA. Association for Computing Machinery.
- 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. Working Paper 27476, National Bureau of Economic Research.
- 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*, 24:427–469.
- OpenAI. 2024a. [Gpt-4.1 prompting guide](#). [https://cookbook.openai.com/examples/gpt4-1\\_prompting\\_guide](https://cookbook.openai.com/examples/gpt4-1_prompting_guide). Accessed: 2025-04-21.
- OpenAI. 2024b. [Hello gpt-4o](#). Accessed: 2025-02-19.
- Mika Saari, A Muzaffar bin Baharudin, Pekka Sillberg, Sami Hyrynsalmi, and Wanglin Yan. 2018. Lora—a survey of recent research trends. In *2018 41st international convention on information and communication technology, electronics and microelectronics (MIPRO)*, pages 0872–0877. IEEE.
- Alexander Scarlatos, Ryan S. Baker, and Andrew Lan. 2025a. Exploring knowledge tracing in tutor-student dialogues using llms. In *Proceedings of the 15th Learning Analytics and Knowledge Conference, LAK 2025, Dublin, Ireland, March 3-7, 2025*. ACM.
- Alexander Scarlatos, Naiming Liu, Jaewook Lee, Richard Baraniuk, and Andrew Lan. 2025b. [Training llm-based tutors to improve student learning outcomes in dialogues](#). *Preprint*, arXiv:2503.06424.
- Shashank Sonkar, Kangqi Ni, Sapana Chaudhary, and Richard Baraniuk. 2024. [Pedagogical alignment of large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13641–13650, Miami, Florida, USA. Association for Computational Linguistics.
- Abhijit Suresh, Jennifer Jacobs, Charis Harty, Margaret Perkoff, James H. Martin, and Tamara Sumner. 2022. The TalkMoves dataset: K-12 mathematics lesson transcripts annotated for teacher and student discursive moves. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4654–4662.
- Anaïs Tack, Ekaterina Kochmar, Zheng Yuan, Serge Bibauw, and Chris Piech. 2023. [The BEA 2023 shared task on generating AI teacher responses in educational dialogues](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 785–795, Toronto, Canada. Association for Computational Linguistics.
- LearnLM Team, Abhinit Modi, Aditya Srikanth Veerubhotla, Aliya Rysbek, Andrea Huber, Brett Wiltshire, Brian Veprek, Daniel Gillick, Daniel Kasenberg, Derek Ahmed, Irina Jurenka, James Cohan, Jennifer She, Julia Wilkowsky, Kaiz Alarakyia, Kevin R. McKee, Lisa Wang, Markus Kunesch, Mike Schaeckermann, Miruna Pîslar, Nikhil Joshi, Parsa Mahmoudieh, Paul Jhun, Sara Wiltberger, Shakir Mohamed, Shashank Agarwal, Shubham Milind Phal, Sun Jae Lee, Theofilos Strinopoulos, Wei-Jen Ko, Amy Wang, Ankit Anand, Avishkar Bhoopchand, Dan Wild, Divya Pandya, Filip Bar, Garth Graham, Holger Winnemoeller, Mahvish Nagda, Prateek Kohar, Renee Schneider, Shaojian Zhu, Stephanie Chan, Steve Yadlowsky, Viknesh Sounderajah, and Yanniss Assael. 2024. [Learnlm: Improving gemini for learning](#). *Preprint*, arXiv:2412.16429.
- Justin Vasselli, Christopher Vasselli, Adam Nohejl, and Taro Watanabe. 2023. [NAISTeacher: A prompt and rerank approach to generating teacher utterances in educational dialogues](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 772–784, Toronto, Canada. Association for Computational Linguistics.

Deliang Wang, Dapeng Shan, Yaqian Zheng, and Gaowei Chen. 2023. Teacher talk moves in k12 mathematics lessons: Automatic identification, prediction explanation, and characteristic exploration. In *Artificial Intelligence in Education*, pages 651–664, Cham. Springer Nature Switzerland.

Rose Wang and Dorottya Demszky. 2024. [Edu-ConvoKit: An open-source library for education conversation data](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 3: System Demonstrations)*, pages 61–69, Mexico City, Mexico. Association for Computational Linguistics.

Rose Wang, Qingyang Zhang, Carly Robinson, Susanna Loeb, and Dorottya Demszky. 2024a. Bridging the novice-expert gap via models of decision-making: A case study on remediating math mistakes. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2174–2199.

Rose E Wang, Ana T Ribeiro, Carly D Robinson, Susanna Loeb, and Dora Demszky. 2024b. Tutor copilot: A human-ai approach for scaling real-time expertise. *arXiv preprint arXiv:2410.03017*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Stella Xin Yin, Zhengyuan Liu, Dion Hoe-Lian Goh, Choon Lang Quek, and Nancy F. Chen. 2025. [Scaling up collaborative dialogue analysis: An ai-driven approach to understanding dialogue patterns in computational thinking education](#). In *Proceedings of the 15th International Learning Analytics and Knowledge Conference, LAK '25*, page 47–57, New York, NY, USA. Association for Computing Machinery.

## A Details on Experimental Setup

For finetuning Llama 3, we perform a hyperparameter search over learning rates [5e-5, 1e-4, 2e-4, 3e-4] and LoRA ranks [4, 8, 16, 32]. For the final training, we set LoRA’s  $\alpha$  to 16, LoRA’s rank to 8, batch size to 64 using gradient accumulation, gradient norm clipping to 1.0, and learning rate to 1e-4. We train for 5 epochs using the AdamW optimizer. We use a random 20% percent of dialogues in the train set to use as a validation set for early stopping. We use the meta-llama/Llama-3.2-3B-Instruct model from the Huggingface Transformers library (Wolf et al., 2019) and run all experiments on NVIDIA L40 GPUs.

We prompt GPT-4o with the OpenAI API using a temperature of 0 while setting the maximum tokens to 1000 and response format to JSON.

Logistic regression is implemented using the sklearn library. The model input is the frequency distribution of moves up to the target tutor move.

For the second-order Markov chain, we compute transition matrices by mapping state pairs based on frequency and normalizing them to yield valid probability distributions.

For the LSTM, we encode the sequence of moves as multi-hot vectors. For multi-label prediction, we use positive weighting for each class, calculated using the proportion of instances in each class. We perform a hyperparameter search with hidden dimensions [64, 128, 256, 512], number of layers [2, 3], dropout rates [0.1, 0.3, 0.5], and learning rates [1e-3, 5e-3, 1e-2]. For multi-label classification, we also search for the optimal probability threshold. Our final models were trained with hidden dimensions 128, 2 layers, dropout of 0.3 and learning rate of 0.001. The learned threshold for multi-label classification is 0.85. We implement the LSTM using Pytorch.

## B Misclassification Analysis

Llama 3 Finetuned on Tutor Move Classification					
Without Previous Tutor Moves			With Previous Tutor Moves		
Ground Truth	Predicted	Count	Ground Truth	Predicted	Count
<b>MathDial</b>					
probing	focus	572	probing	focus	339
telling	focus	143	telling	focus	136
focus	probing	81	focus	probing	124
focus	telling	78	generic	focus	69
generic	focus	46	focus	telling	57
<b>AlgebraNation</b>					
giving_explanation	giving_instruction	73	giving_explanation	giving_instruction	74
giving_instruction	giving_explanation	40	giving_instruction	questioning	29
managing_discussions	asking_for_elaboration	34	asking_for_elaboration	questioning	24
questioning	giving_instruction	34	managing_discussions	questioning	23
giving_instruction	questioning	27	giving_instruction	giving_explanation	23

Table 4: Top 5 misclassifications for MathDial and AlgebraNation with Llama 3, comparing inputs with vs. without previous tutor moves. The most common confusion in MathDial is between *focus* and *probing*. The most common confusion in AlgebraNation is between *giving instruction* and *giving explanation*. Total misclassifications decrease by including previous tutor moves.

GPT-4o on Tutor Move Classification					
Without Previous Tutor Moves			With Previous Tutor Moves		
Ground Truth	Predicted	Count	Ground Truth	Predicted	Count
<b>MathDial</b>					
probing	focus	400	probing	focus	431
focus	probing	274	focus	probing	218
focus	telling	169	focus	telling	180
telling	focus	69	telling	focus	63
telling	probing	60	probing	telling	61
<b>AlgebraNation</b>					
giving_instruction	giving_explanation	95	giving_instruction	giving_explanation	79
giving_explanation	giving_instruction	50	giving_explanation	giving_instruction	44
giving_instruction	correcting	37	encouraging_peer_tutoring	praising_and_encouraging	44
encouraging_peer_tutoring	praising_and_encouraging	37	confirmatory_feedback	praising_and_encouraging	40
confirmatory_feedback	praising_and_encouraging	36	asking_for_elaboration	questioning	38

Table 5: Top 5 misclassifications for MathDial and AlgebraNation with GPT-4o, comparing inputs with vs. without previous tutor moves. The most common confusion in MathDial is between *focus* and *probing*. The most common confusion in AlgebraNation is between *giving instruction* and *giving explanation*.



Llama 3 on Future Tutor Move Prediction					
Without Previous Labels			With Previous Labels		
Ground Truth	Predicted	Count	Ground Truth	Predicted	Count
<b>MathDial</b>					
probing	focus	736	probing	focus	378
telling	focus	335	telling	focus	175
generic	focus	252	generic	focus	125
focus	telling	78	focus	telling	121
probing	telling	64	focus	probing	119
<b>AlgebraNation</b>					
questioning	giving_instruction	105	giving_explanation	giving_instruction	162
giving_explanation	giving_instruction	88	questioning	giving_instruction	131
giving_instruction	giving_explanation	82	confirmatory_feedback	giving_instruction	95
confirmatory_feedback	giving_instruction	65	providing_further_references	giving_instruction	79
managing_discussions	giving_instruction	65	managing_discussions	giving_instruction	78

Table 6: Top 5 misclassifications for MathDial and AlgebraNation with Llama 3, comparing inputs with vs. without previous tutor moves. The most common confusion in MathDial is between focus and probing. Previous labels decrease the total top five misclassifications for MathDial. The most common misclassifications in AlgebraNation all occur when *giving instruction* is predicted.

GPT-4o in Future Tutor Move Prediction					
Without Previous Labels			With Previous Labels		
Ground Truth	Predicted	Count	Ground Truth	Predicted	Count
<b>MathDial</b>					
focus	probing	551	focus	probing	553
probing	focus	311	probing	focus	309
telling	probing	288	telling	probing	283
generic	focus	176	telling	focus	178
telling	focus	173	generic	focus	170
<b>AlgebraNation</b>					
giving_instruction	giving_explanation	144	giving_instruction	giving_explanation	144
giving_explanation	giving_instruction	66	confirmatory_feedback	giving_explanation	62
confirmatory_feedback	giving_explanation	66	questioning	giving_explanation	56
questioning	giving_instruction	65	giving_instruction	confirmatory_feedback	54
questioning	giving_explanation	61	questioning	giving_instruction	54

Table 7: Top 5 misclassifications for MathDial and AlgebraNation with GPT-4o, comparing inputs with vs. without previous tutor moves. The most common confusion in MathDial is between *focus* and *probing*. The most common confusion in AlgebraNation is between *giving instruction* and *giving explanation*. Overall, including previous tutor moves decreases the top 4 misclassifications rates.




Figure 1: Label distribution of tutor move classification for GPT-4o and Llama 3 trained with dialogue and tutor moves.




Figure 2: Label distribution of future tutor move prediction for GPT-4o and Llama 3 trained with dialogue and tutor moves. Llama 3 predictions are heavily skewed towards *giving instruction*.

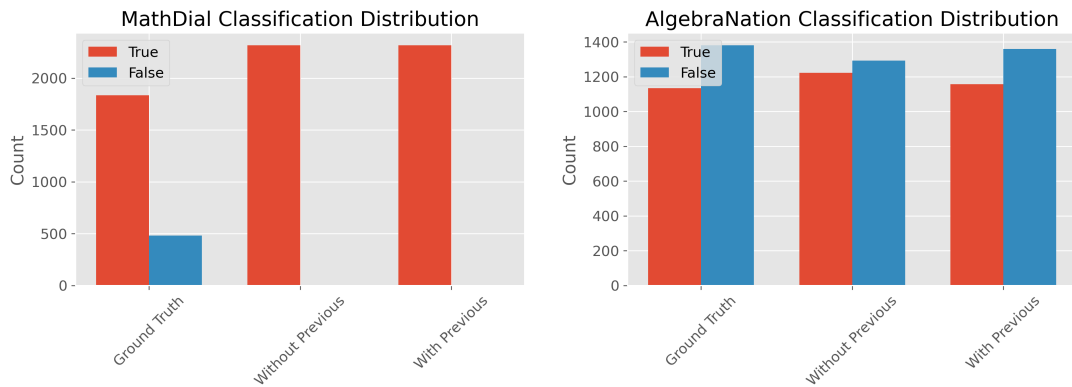


Figure 3: Left: Distribution of dialogue success classification in MathDial using Llama 3. Right: Distribution of dialogue success classification in AlgebraNation using Llama 3.

## C Label Distributions




Figure 4: Tutor move distribution of AlgebraNation dataset. Few classes make up the majority of the distribution.




Figure 5: Tutor move distribution of MathDial dataset.




Figure 6: Left: Distribution of turn-level student success in MathDial. Center: Distribution of final turn student success in MathDial. Right: Distribution of dialogue success in AlgebraNation.

## D Logistic Regression Coefficients and Chi-Squared Analysis

Feature	Coefficient	$\chi^2$	p-value	Significant
confirmatory feedback	0.168682	684.413557	7.330488e-151	Yes
giving instruction	0.144908	505.315750	6.628000e-112	Yes
giving explanation	0.130397	408.828643	6.593282e-91	Yes
praising and encouraging	0.089068	189.507614	4.072012e-43	Yes
giving answers	0.072106	123.889276	8.907816e-29	Yes

Table 8: Logistic regression coefficients and Chi-squared analysis conducted to evaluate the impact of tutor moves on final dialogue correctness, with corresponding p-values for the top 5 significant features for AlgebraNation.

Feature	Coefficient	$\chi^2$	p-value	Significant
generic	1.063463	18.7393	1.4986e-05	Yes
focus	0.175187	2.2684	1.3203e-01	No
probing	-0.186663	6.7967	9.1324e-03	Yes
telling	-1.062664	19.3447	1.0912e-05	Yes

Table 9: Logistic regression coefficients and Chi-squared analysis conducted to evaluate the impact of tutor moves on final dialogue correctness, with corresponding p-values for all features for MathDial.



## **E Prompt Examples for Tutor Move Classification**

### Prompt Used for Classification Task (MathDial)

**System:** You are a math teacher who tutors student on a variety of problems.

**Dialogue ID:** 14

**Task:** Classify the last teacher move into one of the following four categories: *Focus, Probing, Telling, Generic*. These four categories are defined below:

- **Focus**

- Seek Strategy: *So what should you do next?*
- Guiding Student Focus: *Can you calculate ... ?*
- Recall Relevant Information: *Can you reread the question and tell me what is ... ?*

- **Probing**

- Asking for Explanation: *Why do you think you need to add these numbers?*
- Seeking Self Correction: *Are you sure you need to add here?*
- Perturbing the Question: *How would things change if they had ... items instead?*
- Seeking World Knowledge: *How do you calculate the perimeter of a square?*

- **Telling**

- Revealing Strategy: *You need to add ... to ... to get your answer.*
- Revealing Answer: *No, he had ... items.*

- **Generic**

- Greeting/Farewell: *Hi ..., how are you doing with the word problem? Good Job! Is there anything else I can help with?*
- General Inquiry: *Can you walk me through your solution?*

**Instructions:** Given the student-teacher dialogue below, classify the final teacher move. Return a JSON object with the dialogue ID as the key and the teacher move type as the value. For example:

```
{"1234": "focus"}
```

**Dialogue:**

[BEGIN DIALOGUE]

Teacher Turn 0: Hi Ayisha, please talk me through your solution (generic)

Student Turn 0: I started by noting that the first bid was \$200 and each subsequent bid was \$50 more, so the second bid was  $\$200 + \$50 = \$250$ , the third bid was  $\$250 + \$50 = \$300$ , and the fourth bid was  $\$300 + \$50 = \$350$ . Then I realized that Carmen bid after each of the other 3 people, so she made 3 bids in total. Therefore, I concluded that the desk cost her \$350, since that was her winning bid.

Teacher Turn 1: Okay, let's start again. If each bid adds \$50 to the price and there were 3 bids - how much is added to the overall cost? (focus)

Student Turn 1: The overall cost is increased by \$150, since each bid adds \$50 and there were 3 bids.

Teacher Turn 2: Yes \$150. Carmen adds a \$50 bid after each of the three bids. How much more does this add to the cost? (focus)

Student Turn 2: This adds an additional \$150, since each bid is \$50 and there were 3 bids. Therefore, the total cost of the desk is \$350.

Teacher Turn 3:  $3 \times 50$  is 150. So it adds a further \$150 making the total so far \$300 (telling)

Student Turn 3: Yes, that's correct. The total cost of the desk is \$350, since the first bid was \$200 and each subsequent bid was \$50 more, so the final bid was  $\$300 + \$50 = \$350$ .

Teacher Turn 4: No, go back and check what we said. The 3 bids \$150 + Carmen's Bid \$150 + the opening bid \$200 (focus)

Student Turn 4: Yes, you're right. The total cost of the desk is \$350, since the opening bid was \$200, the 3 bids added \$150, and Carmen's bid added an additional \$150, making the total  $\$200 + \$150 + \$150 = \$350$ .

Teacher Turn 5: Check your calculation of  $\$200 + \$150 + \$150 = \$350$ . Your total is not correct

[END DIALOGUE]

Figure 7: Full prompt used to classify the final tutor move in a MathDial dialogue.

### Prompt Used for Classification Task (AlgebraNation)

**System:** You are a math teacher who tutors student on a variety of problems.

**Dialogue ID:** 10520899

**Task:** Classify the last tutor move into one or more of the following categories: *confirmatory\_feedback*, *negative\_feedback*, *correcting*, *giving\_instruction*, *giving\_explanation*, *providing\_further\_references*, *questioning*, *asking\_for\_elaboration*, *praising\_and\_encouraging*, *managing\_frustration*, *managing\_discussions*, *giving\_answers*, *encouraging\_peer\_tutoring*, *guiding\_peer\_tutoring*, *acknowledging\_tutor\_issue*, and *other*. These categories are described below:

- **confirmatory\_feedback:** Whether a reply provides confirmatory feedback about an answer's correctness.
- **negative\_feedback:** Whether a reply states that an answer is incorrect.
- **correcting:** Whether a reply addresses errors in the student's problem-solving approach.
- **giving\_instruction:** Whether a reply breaks down a task, performs a part, or initiates a task for the student to complete.
- **giving\_explanation:** Whether a reply explains concepts, principles, or provides additional information.
- **providing\_further\_references:** Whether a reply includes additional resources or references related to the topic.
- **questioning:** Whether a reply asks questions to stimulate thought or constructive discussion.
- **asking\_for\_elaboration:** Whether a reply requests further details or explanation from the student.
- **praising\_and\_encouraging:** Whether a reply praises or encourages the student for their efforts or successes.
- **managing\_frustration:** Whether a reply addresses the student's negative emotions or frustration.
- **managing\_discussions:** Whether a reply organizes the flow of discussion or adjusts the direction of inquiry.
- **giving\_answers:** Whether a reply directly provides an answer to the posed question.
- **encouraging\_peer\_tutoring:** Whether a reply promotes tutoring interactions among peers.
- **guiding\_peer\_tutoring:** Whether a reply provides feedback on peer tutoring interactions.
- **acknowledging\_tutor\_issue:** Whether the tutor expresses uncertainty in their reply.
- **other:** Binary indicator for tutoring strategies not classified under the existing labels.

**Instructions:** Given the student-teacher dialogue below, classify the final teacher move. Return a JSON object with the dialogue ID as the key and the teacher move type(s) as the value. For example:

```
{"1234": ["confirmatory_feedback", "correcting"]}
```

**Dialogue:**

[BEGIN DIALOGUE]

Student: Can someone help me?

Student: You have to plug in zeros for x and y right

Tutor: Get it into  $y=mx+b$  form. ['giving\_instruction']

Tutor: SO bring  $6x$  to the right side first. ['giving\_instruction']

Teacher: Okay Patrice, you want to put that in slope-intercept form ['giving\_instruction']

Student:  $-5y=30+6x$

Teacher: Now isolate y ['giving\_instruction']

Student:  $-5$  on both sides ?

Student: or divide

Teacher: You would divide

[END DIALOGUE]

Figure 8: Full prompt used to classify the final tutor move in an AlgebraNation dialogue.