



Multi-granularity Time-based Transformer for Knowledge Tracing

Tong Zhou

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

April 9, 2023

Multi-granularity Time-based Transformer for Knowledge Tracing

Tong Zhou

Department of Computer Science

Johns Hopkins University

Baltimore, United States

tzhou11@jhu.edu

Abstract—In this paper, we present a transformer architecture for predicting student performance on standardized tests. Specifically, we leverage students’ historical data, including their past test scores, study habits, and other relevant information, to create a personalized model for each student. We then use these models to predict their future performance on a given test. Applying this model to the RIID dataset, we demonstrate that using multiple granularities for temporal features as the decoder input significantly improve model performance. Our results also show the effectiveness of our approach, with substantial improvements over the LightGBM method. Our work contributes to the growing field of AI in education, providing a scalable and accurate tool for predicting student outcomes.

Index Terms—Transformer, Multi-granularity, Education, RIID, Deep Learning

I. INTRODUCTION

Knowledge tracing is an important field of research in educational data mining, as it can help to improve the effectiveness and efficiency of learning. The first application is personalized learning. [1] surveys that by tracking the histories of individual students over time, instructors can tailor instructional materials to meet specific needs of each student. The second application is early intervention. [2] recognizes that knowledge tracing enables instructors to intervene early and offer extra assistance or resources to students who need it in order to succeed by identifying those who are having difficulty with specific concepts or skills. Knowledge tracing can also facilitate adaptive learning, which can change the level of difficulty and pace of instruction based on the understanding and performance of individual students. [3] shows that learning and engagement can be enhanced since students are appropriately challenged and can avoid frustration or boredom.

Recent advances in artificial intelligence (AI) have shown great promise for improving educational outcomes. In particular, deep learning models have been developed to predict student performance on standardized tests, such as the SAT and TOEIC, based on a variety of factors, including their previous academic records, socio-economic background, and personal characteristics. One of the key challenges in building such models is capturing the dynamic and complex nature of student behavior, which can vary widely over time and across individuals.

In this paper, we propose a Transformer that leverages users’ histories for educational performance prediction. The

Transformer is a state-of-the-art neural network architecture that has achieved impressive results in various natural language processing tasks, such as machine translation, question answering, and text generation. Our approach extends the Transformer to model student performance by incorporating their past academic records, study habits, and other relevant information.

To evaluate the effectiveness of our approach, we conducted experiments on real-world educational datasets, including the Kaggle Riid AIED Challenge dataset. We demonstrate that converting temporal features into multiple categorical features with different granularities can greatly improve model’s performance. The results even show that the lecture information is irrelevant in the presence of the multi-granularity temporal features. Our results also demonstrate that our model outperforms traditional LightGBM, achieving state-of-the-art performance in predicting student performance on standardized tests. Moreover, we show that our model can be used to provide personalized recommendations to students based on their historical data, enabling them to improve their academic performance.

II. RELATED WORK

Works of knowledge tracing models mainly follow two different approaches: BKT (Bayesian Knowledge Tracing) and DKT (Deep Knowledge Tracing). BKT is a probabilistic model where student knowledge is modeled as a latent variable and other observed context information and learning performance are used to identify the latent structure represented by a Hidden Markov Model [2, 4, 5, 6]. BKT, however, has some innate limitations, such as its implausible assumption of independence between skills, and its inability to address correlations between different skills or knowledge components.

The Deep Knowledge Tracing (DKT) has received increasing attention. DKT relies on a recurrent neural network (RNN) architecture, using a Long Short-Term Memory (LSTM) network to take a sequence of student responses and other context information as input. By its design, DKT is able to handle correlations between different knowledge components and capture complex interactions between these components over time [7, 8, 9].

Using Transformer architecture has gained more popularity in recent years, as its self attention mechanism has demon-




Fig. 1. Model Architecture

strated great effectiveness for sequential prediction tasks. Three notable papers based on Transformer are SAKT [10], SAINT [11] and SAINT+ [12]. SAINT+ is mostly related to our paper. We present a slightly different neural network architecture and a notable improvement in feature engineering for temporal features.

III. METHODS

We present our Transformer architecture in Figure 1. The encoder in our model consists of four components: question embeddings, part embeddings, position embeddings and prior-question-had-explanation embeddings. The decoder in our model consists of 6 components: position embeddings, response embeddings, prior-elapsed-time embeddings, lag-time-1 categorical embeddings, lag-time-2 categorical embeddings, lag-time-3 categorical embeddings. One of the notable features of our Transformer architecture is that we included three separate embedding layers for representing lag time in three different temporal granularities: by seconds, by minutes, and by days.

This architecture captures multiple types of information. The encoder captures information about the questions and parts of the questions, as well as the presence or absence of explanations. This allows the model to capture multiple types of information that can affect student learning, and to learn meaningful representations of these factors.

It also incorporates temporal information. The encoder uses position embedding to encode the position of each token within the sequence of interactions between the student and the instruction materials. This allows the model to capture the temporal dynamics of the data and learn representations that can account for changes in student learning over time.

It can model complex interactions: The use of multiple embedding layers in both the encoder and decoder components of the model allows for the modeling of complex interactions between different factors that can affect student learning. For example, the model can learn to capture the relationship between a student’s prior response and their subsequent performance on a related question, or the interaction between elapsed time and lag time in predicting student learning outcomes.

A. Data

We use the dataset provided by Kaggle from the “Riiid! Answer Correctness Prediction” competition¹, which challenges participants to build machine learning models that can predict students’ responses to questions on an educational platform.

The dataset contains over 100 million rows of data from student interactions with the platform, including information on the questions, answers, and explanations provided, as well as metadata such as the time elapsed since the previous interaction and the student’s performance history. The goal is to predict whether or not a student will correctly answer the next question in a sequence, based on their past interactions with the platform.

B. Encoder

- 1) **Question embeddings:** this feature encodes question ID. We used an embedding layer of 128 dimensions. The question embedding layer is intended to help the neural network identify the connections between various abilities or concepts and their significance to the student’s overall learning development. The network may learn to spot links between question difficulty and student performance by modeling the questions and tasks using an embedding layer.
- 2) **Part embeddings:** this feature encodes different sections of test in dataset. This embedding layer is to help the neural network understand the connections between different parts of the questions and their significance to the student’s overall learning progress. The network may learn to capture the connections between category difficulty and student performance by encoding the categories using an embedding layer.
- 3) **Position embeddings:** this feature serves to encode the relative position of each token in the input sequence, since the self-attention mechanism treats all tokens equally regardless of their position in the sequence.
- 4) **Prior-question-had-explanation embeddings:** this feature encodes whether or not a user saw the explanation of a question after answering it, with a vocabulary

¹<https://www.kaggle.com/c/riiid-test-answer-prediction/>

size of 3. This embedding layer is intended to help the neural network identify the connection between the user’s attempt to see an explanation and their overall learning outcomes.

C. Decoder

- 1) **Position embeddings:** this feature serves to encode the relative position of each token in the input sequence, since the self-attention mechanism treats all tokens equally regardless of their position in the sequence.
- 2) **response embeddings:** this feature encodes if a user’s answer is correct or not, with a vocabulary size of 3. The goal of the response embedding layer is to help the neural network understand the connection between the accuracy of the student’s response and their overall progress of learning. The network can learn to distinguish between right and wrong responses can record patterns and links between response accuracy and student performance by modelling the response using an embedding layer. This can increase the network’s predictive accuracy and make it easier to spot areas where a student may be struggling or excelling.
- 3) **prior-elapsed-time embeddings:** this feature encodes the amount of time that a user used to answer the question, with a vocabulary size of 301. The goal of the elapsed time embedding is to assist the neural network in learning meaningful representations of the temporal information and in capturing the amount of time that has passed between the current task or question and the previous one. The network may learn to distinguish between various time periods and capture patterns and correlations between elapsed time and student performance by expressing elapsed time using an embedding layer. This can increase the network’s prediction accuracy and make it easier to spot long-term trends in student learning.
- 4) **lag-time-1 categorical embeddings:** this feature encodes lag time in seconds, with a vocabulary size of 301.
- 5) **lag-time-2 categorical embeddings:** this feature encodes lag time for representing lag time in minutes, with a vocabulary size of 1441.
- 6) **lag-time-3 categorical embedding:** this feature encodes lag time for representing lag time in days, with a vocabulary size of 366.

D. Training

Binary cross-entropy loss is used to evaluate our model. Since the dataset contains almost 100 million rows of data, we simply use 97.5% of data as the training data, and 2.5% as the validation data. The ultimate hyper-parameters are as follows

- Max sequence: 100
- Embedding size: 128
- Number of layers in encoder: 2
- Number of layers in decoder: 2
- Batch size: 256

- Dropout: 0.1
- Epoch: 10
- Number of heads: 8
- Optimizer: AdamW
- Learning rate: $5e - 4$

IV. RESULTS

We did three experiments for comparison purposes: Transformer without lecture information, Transformer with lecture information, and LightGBM

TABLE I
COMPARISON OF AUC BETWEEN TRANSFORMER VARIANTS AND LIGHTGBM

Model	AUC
Transformer with lecture	0.800
Transformer without lecture	0.800
LightGBM	0.785

Table I shows that the multiple temporal granularities are significant such that even lecture information is ignored, they can capture the complex interactions between question difficulty and student performance.

V. CONCLUSION

In this paper, we introduced a novel Transformer architecture for knowledge tracing, which is designed to capture multiple types of information that can affect student learning. Our architecture includes a sophisticated encoder component that captures information about the questions, parts of questions, and presence or absence of explanations, as well as a decoder component that incorporates temporal information about the sequence of interactions between the student and the learning materials. Notably, we included three separate embedding layers to represent lag time in different temporal granularities, allowing the model to capture the temporal dynamics of the data.

Our architecture also has the ability to model complex interactions between different factors that can affect student learning, including the relationship between a student’s prior response and their subsequent performance on a related question, or the interaction between elapsed time and lag time in predicting student learning outcomes. Overall, our Transformer architecture provides a powerful tool for accurately predicting student performance and identifying areas where additional support or guidance may be needed.

ACKNOWLEDGMENT

The author would like to thank Kaggle and Riiid AIED for organizing the competition and providing such high quality data.

REFERENCES

- [1] G. Mahajan and B. Saini, “Educational data mining: A state-of-the-art survey on tools and techniques used in edm,” *International Journal of Computer Applications*

- & *Information Technology*, vol. 12, no. 1, pp. 310–316, 2020.
- [2] A. T. Corbett and J. R. Anderson, “Knowledge tracing: Modeling the acquisition of procedural knowledge,” *User modeling and user-adapted interaction*, vol. 4, pp. 253–278, 1994.
 - [3] N. T. Heffernan and C. L. Heffernan, “The assistments ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching,” *International Journal of Artificial Intelligence in Education*, vol. 24, pp. 470–497, 2014.
 - [4] R. S. d. Baker, A. T. Corbett, and V. Aleven, “More accurate student modeling through contextual estimation of slip and guess probabilities in bayesian knowledge tracing,” in *Intelligent Tutoring Systems: 9th International Conference, ITS 2008, Montreal, Canada, June 23-27, 2008 Proceedings 9*. Springer, 2008, pp. 406–415.
 - [5] J. Kasurinen and U. Nikula, “Estimating programming knowledge with bayesian knowledge tracing,” *ACM SIGCSE Bulletin*, vol. 41, no. 3, pp. 313–317, 2009.
 - [6] M. V. Yudelson, K. R. Koedinger, and G. J. Gordon, “Individualized bayesian knowledge tracing models,” in *Artificial Intelligence in Education: 16th International Conference, AIED 2013, Memphis, TN, USA, July 9-13, 2013. Proceedings 16*. Springer, 2013, pp. 171–180.
 - [7] C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein, “Deep knowledge tracing,” *Advances in neural information processing systems*, vol. 28, 2015.
 - [8] J. Zhang, X. Shi, I. King, and D.-Y. Yeung, “Dynamic key-value memory networks for knowledge tracing,” in *Proceedings of the 26th international conference on World Wide Web*, 2017, pp. 765–774.
 - [9] Q. Liu, Z. Huang, Y. Yin, E. Chen, H. Xiong, Y. Su, and G. Hu, “Ekt: Exercise-aware knowledge tracing for student performance prediction,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 33, no. 1, pp. 100–115, 2019.
 - [10] S. Pandey and G. Karypis, “A self-attentive model for knowledge tracing,” *arXiv preprint arXiv:1907.06837*, 2019.
 - [11] Y. Choi, Y. Lee, D. Shin, J. Cho, S. Park, S. Lee, J. Baek, C. Bae, B. Kim, and J. Heo, “Ednet: A large-scale hierarchical dataset in education,” in *Artificial Intelligence in Education: 21st International Conference, AIED 2020, Ifrane, Morocco, July 6–10, 2020, Proceedings, Part II 21*. Springer, 2020, pp. 69–73.
 - [12] D. Shin, Y. Shim, H. Yu, S. Lee, B. Kim, and Y. Choi, “Saint+: Integrating temporal features for ednet correctness prediction,” in *LAK21: 11th International Learning Analytics and Knowledge Conference*, 2021, pp. 490–496.