



Option Tracing: Beyond Correctness Analysis in Knowledge Tracing

Aritra Ghosh¹, Jay Raspat², and Andrew Lan¹(✉)

¹ University of Massachusetts Amherst, Amherst, USA

andrewlan@cs.umass.edu

² Pittsburgh, PA, USA

Abstract. Knowledge tracing refers to a family of methods that estimate each student’s knowledge component/skill mastery level from their past responses to questions. One key limitation of most existing knowledge tracing methods is that they can only estimate an *overall* knowledge level of a student per knowledge component/skill since they analyze only the (usually binary-valued) correctness of student responses. Therefore, it is hard to use them to diagnose specific student errors. In this paper, we extend existing knowledge tracing methods beyond correctness prediction to the task of predicting the exact option students select in multiple choice questions. We quantitatively evaluate the performance of our option tracing methods on two large-scale student response datasets. We also qualitatively evaluate their ability in identifying common student errors in the form of clusters of incorrect options across different questions that correspond to the same error.

1 Introduction

Knowledge tracing (KT) [9] refers to a family of student modeling methods that estimate student mastery levels on knowledge components/skills/concepts from their past responses to questions/items and predict their future performance. These estimates and predictions can be used to i) provide feedback to students on their progress, especially in intelligent tutoring systems [44] and ii) drive personalization, i.e., selecting the action that each learner should take next to maximize their learning outcomes [10, 28, 36]. Many different KT methods have been developed, ranging from hidden Markov model-based Bayesian knowledge tracing methods [21, 33, 46], factor analysis-based methods such as learning factor analysis [5], performance factor analysis [34], and the item Difficulty, student ability, skill, and student skill practice history (DAS3H) method [7], to deep learning-based methods [15, 31, 32, 35, 45, 47]. These methods have enjoyed various degrees of success; some of these methods, including most Bayesian knowledge tracing and factor analysis-based methods, exhibit excellent interpretability while other, deep learning-based methods trade off interpretability for excellent predictive accuracy on students’ future performance.

This work is supported by the National Science Foundation under grant IIS-1917713. We also thank the reviewers for their constructive feedback.

© Springer Nature Switzerland AG 2021

I. Roll et al. (Eds.): AIED 2021, LNAI 12748, pp. 1–13, 2021.

https://doi.org/10.1007/978-3-030-78292-4_12




Fig. 1. Some distractor options in well-designed MCQs are potentially capable of capturing typical student errors. Option C in both questions here correspond to the error of not mastering the order of operations and always working left to right.

However, one key limitation of these KT methods is that they operate exclusively on (usually binary-valued) response data that indicates whether a student responds to a question correctly or not. Therefore, they can only estimate students’ *overall* mastery level on each knowledge component. However, not all incorrect responses are equal: there can be numerous incorrect ways to answer a math question [27], caused by different underlying errors. Studies have shown that only a fraction of incorrect answers generated by students can be anticipated and explained by cognitive models integrated into intelligent tutoring systems [24, 36, 41], teachers [11], and numerical simulations [11, 37]. Typical underlying errors include having a “buggy rule” [4], exhibiting a certain misconception [12, 13, 38], or a general lack of knowledge on certain knowledge components [2]. Since it is hard to diagnose such student errors from correctness data alone, we need to develop KT methods that analyze full student responses.

Some datasets, including the large-scale Eedi¹ [43] and EdNet² [8] datasets, contain the exact options students select on multiple choice questions (MCQs); this option data provides us with an opportunity to extend existing KT methods to analyze specific student option selections rather than their answer correctness. In an ideal situation, well-designed MCQs should have well-crafted incorrect distractor options that each corresponds to one or more typical student errors; Figure 1 shows an example from the Eedi dataset for two questions on the subject brackets, indices, division, multiplication, addition, subtraction (BIDMAS). Option C in both questions correspond to the same error of not fully mastering “order of operations” and always working left to right. However, manually identifying these errors is an unscalable and labor-intensive process since most existing MCQs do not come with consistent labels on the error(s) underlying each incorrect option. Therefore, it is important to explore whether we can develop KT methods to identify errors each incorrect option corresponds to and potentially diagnose student errors automatically. These methods would then be useful through i) informing teachers to communicate with students to understand the source of their errors, ii) enabling the development of automated

¹ <https://eedi.com/projects/neurips-education-challenge>.

² <https://github.com/rriid/ednet>.

feedback [19], and iii) enabling the design of alternative instructional approaches such as asking students to criticize erroneous examples [1].

1.1 Contributions

In this paper, we develop option tracing (OT), a KT framework that uses the exact option each student selects on each question as both input and predicted output. We extend several existing KT methods to the OT setting, including a long short-term memory (LSTM) network-based method, deep knowledge tracing (DKT) [35], a graph convolutional network-based method, graph-based interaction model for knowledge tracing (GIKT) [45], and an attention network-based method, attentive knowledge tracing (AKT) [15]. We emphasize that the goal of this paper is **NOT** to compare all KT methods; instead, our goal is to study how can we generalize them to analyze student option selections in MCQs. Therefore, we only study some representative methods. We conduct the following experiments on the Eedi and EdNet datasets: First, we quantitatively evaluate our OT methods under both the collaborative filtering (CF) setup (introduced by the NeurIPS 2020 Education Challenge [43]) and the typical KT setup on the task of option prediction. Second, we qualitatively demonstrate the interpretability exhibited by our OT framework using clustering algorithms to group incorrect options across multiple questions into clusters of shared underlying errors. Results show that the learned clusters match up with those manually identified by a domain expert to some degree. Therefore, OT can potentially offer a *bottom-up* approach for error identification by extracting student errors from actual data instead of the typical *top-down* approach of anticipating errors before seeing data. Our implementation will be publicly available at <https://github.com/arghosh/OptionTracing>.

2 Related Work

The options students select in MCQs can be regarded as a type of *categorical data*, which has previously been studied in both the item response theory (IRT) and recommender systems research communities. However, in both cases, most prior works focus on the case where the categories are *ordered*. In IRT research, polytomous IRT-based models [25, 26, 30] are used to model students’ responses with multiple ordered categories, such as letter grades and partial credits. In recommender systems research, neural collaborative filtering (NCF)-based methods are used to model star ratings provided by users on items [20]. There are relatively few models for *unordered* categorical data such as the nominal response model (NRM) from the IRT research community, which has been applied to the analysis of MCQs [40, 42].

3 Data and Problem Setup

The Eedi dataset contains the responses of more than 100,000 students to 27,613 MCQs across 389 labeled subjects, totaling over 15 million responses over the course of more than a year. Each response corresponds to the exact option a student selected on each question (among four options, $\{A,B,C,D\}$). We will also use a small subset of the Eedi data where we have access to the exact question (in the form of images) for quantitative analysis; this dataset contains the responses of more than 4,900 students to 948 questions, totaling in over 1.3 million responses. The EdNet dataset contains the responses of more than 700,000 students to 13,169 MCQs across 189 labeled subjects, totaling over 95 million responses over the course of more than two years.

We use two experimental setups for evaluation purposes. First, in the CF setup, the task is to predict each student’s responses to a subset of questions that they responded to, given their responses to other questions (possibly in the future). Popular methods for this setup are neural collaborative filtering (NCF) [20] and graph convolutional networks (GCN) [3, 23]. Second, in the KT setup for evaluating KT methods, the task is to predict each student’s responses to future questions based on their entire past response history.

3.1 Problem Setup

Each student’s performance record consists of a sequence of responses to questions assigned at a series of discrete time steps. For student i at time step t , we denote the combination of the question that they answered, the set of subjects this question covers, their binary-valued response correctness, the option they chose, and the correct option to this question as a tuple, $(q_t^i, \{s_{t,j}^i\}_{j=1}^{n_t^i}, r_t^i, y_t^i, c_t^i)$, where $q_t^i \in \mathbb{N}^+$ is the question index, $s_{t,j}^i \in \mathbb{N}^+$ denotes the index of the j^{th} subject, $j \in 1, \dots, n_t^i$ since each question can be tagged with multiple subjects, $r_t^i \in \{0, 1\}$ is the response correctness (1 corresponds to a correct response), $y_t^i \in \{A, B, C, D\}$ is the option the student selected, and $c_t^i \in \{A, B, C, D\}$ is the correct option for this question. In the CF setup, we associate a mask variable $m_t^i \in \{0, 1\}$ with each time step, where 1 represents that the timestep is part of the training set. This variable helps us to mask out responses we need to predict when we compute the training loss. Given observed responses $\{(q_t^i, \{s_{t,j}^i\}_{j=1}^{n_t^i}, r_t^i, y_t^i, c_t^i)\}_{t:m_t^i=1}$, the task is to predict the exact options students select on questions in the test set, i.e., $y_{t'}^i$ for $(t', i') : m_{t'}^i = 0$. In the KT setup, we observe each student’s entire history of responses to questions; thus, given their past history up to time $t-1$ as $\{(q_\tau^i, \{s_{\tau,j}^i\}_{j=1}^{n_\tau^i}, r_\tau^i, y_\tau^i, c_\tau^i)\}_{\tau=1}^{t-1}$, our goal is to predict y_t^i at the current time step, t . Under these notations, existing KT methods focus on predicting response correctness, r_t^i .

4 Methodology

In this section, we detail our OT methods for both the CF and KT setups. Before delving into the individual methods, we start with a set of unified modules that apply to all methods in this paper. The question embedding module $E_q : q \rightarrow \mathbb{R}^d$ transforms the question index q_t^i to a d -dimensional, learnable real-valued vector in \mathbb{R}^d . Similarly, the response embedding module $E_r : r \rightarrow \mathbb{R}^d$ transforms the response correctness r_t^i to \mathbb{R}^d and the option embedding module $E_o : \{A, B, C, D\} \rightarrow \mathbb{R}^d$ transforms the correct option c_t^i and the chosen option y_t^i to vectors in \mathbb{R}^d . We do not use separate embeddings for every question-option (q, o) pair since that leads to overfitting in our experiments; instead, the $2d$ -dimensional embedding for (q, o) is obtained using $[E_q(q) \oplus E_o(o)]$ where \oplus is the concatenation operator. The subject embedding module $E_s : s \rightarrow \mathbb{R}^d$ transforms the subject index to \mathbb{R}^d . Since each question may be tagged with several subjects, we define the final subject embedding as $E_s(\{s_{t,j}^i\}_{j=1}^{n_i}) = \sum_{j=1}^{n_i} E_s(s_{t,j}^i)$. Some of the methods (such as NCF) use a user embedding module $E_u : i \rightarrow \mathbb{R}^d$ that transforms the student index to \mathbb{R}^d . For simplicity, we use the same d -dimensional vector for all embedding modules; however, the dimensions of each module can be different. We train all model parameters, denoted as Θ , which contains the embeddings listed here and other model parameters specific to each individual method, by minimizing the negative log-likelihood of the selected options as

$$\underset{\Theta}{\text{minimize}} \quad - \sum_{i=1}^{|\text{Students}|} \sum_{t=1}^{|\text{Sequence}_i|} \sum_{o \in \{A, B, C, D\}} \mathbb{1}[y_t^i = o] \log p(o|q_t^i; \Theta),$$

where $\mathbb{1}$ is the indicator function. Since the options are unordered categories, the resulting loss function corresponds to the common cross-entropy loss [16].

4.1 Option Prediction Under the CF Setup

NCF. NCF is one of the most popular CF methods for user-item interaction data. In the option prediction task, students correspond to users and questions corresponds to items. The input for NCF at time step t for student i , \mathbf{x}_t^i , is

$$\mathbf{x}_t^i = [E_q(q_t^i) \oplus E_u(i) \oplus E_s(\{s_t^i\})].$$

Predictive probabilities $p(y_t^i = o)$ over four options $o \in \{A, B, C, D\}$ are calculated using the softmax function [16],

$$\mathbf{z}_t^i = f(\mathbf{x}_t^i) \in \mathbb{R}^4, \quad p(y_t^i = o|\mathbf{x}_t^i) = [\text{softmax}(\mathbf{z}_t^i)]_o, \quad \hat{y}_t^i = \underset{o \in \{A, B, C, D\}}{\text{argmax}} [\mathbf{z}_t^i]_o,$$

where $f(\cdot)$ denotes a feed-forward, fully-connected neural network and $[\cdot]_o$ refers to the o^{th} entry of a vector. In NCF, the model parameters are the weights and biases in the feed-forward neural network $f(\cdot)$; this prediction module is shared by the subsequent methods.

PO-BiDKT. The main drawback of NCF is that the student embedding is static and not updated as students answer more questions and their knowledge states evolve. Recurrent neural networks, and in particular LSTM-type models are capable of modeling evolving knowledge as hidden states [35]. However, we cannot directly use methods such as DKT in the CF setup since the student’s responses at some time steps in their response sequence are not observed. Therefore, we use the following method to handle evolving knowledge states using recurrent networks with missing observations. The input at each time step is given by

$$\mathbf{x}_t = [\mathbf{E}_q(q_t^i) \oplus \mathbf{E}_o(c_t^i) \oplus \mathbf{E}_s(\{s_t^i\}) \oplus (\mathbf{E}_o(y_t^i) \odot m_t^i) \oplus (\mathbf{E}_l(r_t^i) \odot m_t^i)], \quad (1)$$

where \odot denotes the element-wise multiplication between two vectors. We mask the option embeddings and response correctness embeddings using m_t^i for time steps where we do not observe them but still use the question embedding as input. We also extend the base LSTM module in DKT to a bi-directional LSTM (Bi-LSTM) [17]. Here, we compute two latent knowledge states using two separate LSTM modules, the forward state $\vec{\mathbf{h}}_t$ that summarizes the student’s past response history and the backward state $\overleftarrow{\mathbf{h}}_t$ that summarizes the student’s future response history at time step t as

$$\vec{\mathbf{h}}_{t+1} = \text{Forward LSTM}(\vec{\mathbf{h}}_t, \mathbf{x}_t), \quad \overleftarrow{\mathbf{h}}_{t-1} = \text{Backward LSTM}(\overleftarrow{\mathbf{h}}_t, \mathbf{x}_t).$$

The final latent knowledge state is the concatenation of the two states as $\mathbf{h}_t = [\vec{\mathbf{h}}_t \oplus \overleftarrow{\mathbf{h}}_t]$. The parameters include two sets of parameters for the forward and backward LSTMs in addition to the parameters for the fully connected network $f(\cdot)$. We call this method partially observed bi-directional DKT, or PO-BiDKT. The output to the prediction module is computed using

$$\mathbf{z}_t^i = f([\mathbf{h}_t^i \oplus \mathbf{E}_q(q_t^i) \oplus \mathbf{E}_o(c_t^i) \oplus \mathbf{E}_s(\{s_t^i\})]) \in \mathbb{R}^4. \quad (2)$$

GCN-Augmented PO-BiDKT (BiGIKT). In our datasets, each question is tagged with a few subjects by question designers or domain experts. These subject tags provide important information on how these questions are related since we expect questions from the same subject to have some shared features. GCNs excel at formulating these relations and learning from graph-structured data. Since we can represent the question-subject association matrix using a bipartite graph, (loosely) following GIKT [45], we connect GCNs with PO-BiDKT to jointly learn question and subject embeddings using the structure imposed by the subject tags. In this method, we use hierarchical representations of subjects and questions: starting with initial subject and question embeddings $\mathbf{E}_s(s_j^i)$ and $\mathbf{E}_q(q_i^i)$, the first layer GCN embedding for the j^{th} subject and the second layer GCN embedding for the i^{th} question are computed as

$$\mathbf{s}_j^1 = \tanh\left(\mathbf{W}_s^s \mathbf{E}_s(s_j) + \frac{\sum_{i \in N_j^s} \mathbf{W}_s^q \mathbf{E}_q(q_i)}{|N_j^s|}\right), \quad \mathbf{q}_i^2 = \tanh\left(\mathbf{W}_q^q \mathbf{E}_q(q_i) + \frac{\sum_{j \in N_i^q} \mathbf{W}_q^s \mathbf{s}_j^1}{|N_i^q|}\right),$$

where $N_j^s(N_i^q)$ denotes the set of questions (subjects) associated with subject (question) $s_j(q_i)$ and $\mathbf{W}_s^s, \mathbf{W}_s^q, \mathbf{W}_q^s$ and \mathbf{W}_q^q are learnable parameter matrices. The hyperbolic tangent (tanh) non-linearity operate entry-wise on vectors. We replace the subject embeddings $\mathbf{E}_s(s_t^i)$ and the question embedding $\mathbf{E}_q(q_t^i)$ in the base Bi-LSTM (Eq. 1 and Eq. 2) with these GCN-based embeddings. The model parameters of this method include the GCN weight parameter matrices in addition to the Bi-LSTM parameters.

4.2 Option Prediction Under the KT Setup

In the KT setup, we predict future responses using only past responses and assume that every past student response is observed. We extend several existing neural network-based KT methods for the option prediction task.

DKT. We apply a simple modification to the DKT method [35] to extend it to i) predict options instead of response correctness and ii) handle questions that are tagged with multiple subjects (the original DKT method assumes that each question is tagged with a single subject). We use

$$\mathbf{x}_t = [\mathbf{E}_q(q_t^i) \oplus \mathbf{E}_o(c_t^i) \oplus \mathbf{E}_s(\{s_t^i\}) \oplus \mathbf{E}_o(a_t^i) \oplus \mathbf{E}_l(r_t^i)]$$

as the input to the DKT LSTM input module. The student’s hidden knowledge states are computed using the LSTM model as $\mathbf{h}_{t+1} = \text{LSTM}(\mathbf{h}_t, \mathbf{x}_t)$. The predictive probabilities of selecting each option are computed using

$$\mathbf{z}_t^i = f([\mathbf{h}_t^i \oplus \mathbf{E}_q(q_t^i) \oplus \mathbf{E}_o(c_t^i) \oplus \mathbf{E}_s(\{s_t^i\})]) \in \mathbb{R}^4, \hat{y}_t^i = \underset{o \in \{A, B, C, D\}}{\operatorname{argmax}} [\mathbf{z}_t^i]_o.$$

DKVMN. Instead of using LSTMs to model latent knowledge state transitions, the dynamic key-value memory network (DKVMN) method uses a key-value memory network to retrieve and update knowledge at every time step using an external memory module as $\mathbf{h}_{t+1} = \text{MemoryModule}(\mathbf{h}_t, \mathbf{x}_t)$; refer to [47] for details. We use the same input and output structure for the DKVMN memory module as that for DKT.

AKT. We also adapt AKT, an attention network-based, state-of-the-art KT method for the option prediction task. AKT computes a query, a key, and a value vector for each time step, and then uses the similarity between the query and key vectors at different time steps to attend to questions in the past and use their corresponding value vectors to retrieve acquired knowledge in the past. We compute the query, key, and value vectors as $\mathbf{q}_t = \mathbf{W}^Q \mathbf{n}$, $\mathbf{k}_t = \mathbf{W}^K \mathbf{n}$, and $\mathbf{v}_t = \mathbf{W}^V [\mathbf{E}_l(r_t^i) \oplus \mathbf{E}_q(q_t^i) \oplus \mathbf{E}_o(y_t^i) \oplus \mathbf{E}_o(c_t^i)]$ respectively, where \mathbf{W}^Q , \mathbf{W}^K , and \mathbf{W}^V are the query, key, and value projection matrices and $\mathbf{n} = [\mathbf{E}_q(q_t^i) \oplus \mathbf{E}_s(\{s_t^i\}) \oplus \mathbf{E}_o(c_t^i)]$. The retrieved latent knowledge state is then computed as $\mathbf{h}_t = g\left(\sum_{\tau < t} \alpha_{t,\tau} \mathbf{v}_\tau\right)$, where g is another feedforward network and $\alpha_{t,\tau}$ is the normalized attention score between the query at the current time step t and the key at a past time step τ . For AKT, we employ the exponential decay module to compute the attention scores [15] and then compute the output using the attention-weighted value \mathbf{h}_t^i and a fully connected network $f(\cdot)$.

Table 1. Performance of all methods under the CF (top half) and KT (bottom half) setups on both datasets. Best results are in **bold**.

Model	Option prediction				Correctness prediction	
	Accuracy		Average Macro F_1 Score		Accuracy	
	Eedi	EdNet	Eedi	EdNet	Eedi	EdNet
NCF	64.75 \pm 0.02	67.24 \pm 0.01	0.2824 \pm 0.002	0.2552 \pm 0.001	72.6 \pm 0.03	71.49 \pm 0.01
PO-BiDKT	65.87 \pm 0.01	69.42 \pm 0.01	0.3283 \pm 0.001	0.3260 \pm 0.001	75.18 \pm 0.01	75.21 \pm 0.02
BiGIKT	66.16 \pm 0.02	69.29 \pm 0.02	0.3261 \pm 0.001	0.3168 \pm 0.001	75.62 \pm 0.02	75.07 \pm 0.01
DKT	65.95 \pm 0.44	68.03 \pm 0.09	0.313 \pm 0.008	0.2887 \pm 0.005	74.7 \pm 0.34	73.19 \pm 0.06
DKVMN	66.03 \pm 0.49	68.01 \pm 0.1	0.3152 \pm 0.007	0.2842 \pm 0.005	74.75 \pm 0.3	73.02 \pm 0.06
AKT	65.91 \pm 0.47	68.44 \pm 0.09	0.3139 \pm 0.007	0.3062 \pm 0.004	74.65 \pm 0.31	73.6 \pm 0.06

5 Experiments

Experimental Setup. In addition to the option prediction task, we also evaluate all methods under the standard, binary-valued response correctness prediction task. We do not need to use a separate set of methods; instead, we can simply replace the final output layer of the option predictor module ($f : \cdot \rightarrow \mathbb{R}^4$) with an output layer that consists of a single node ($f : \cdot \rightarrow \mathbb{R}^1$) for all OT methods; the resulting loss function corresponds to standard binary cross entropy loss. For option prediction, we use both accuracy and macro F_1 score as evaluation metrics. For correctness prediction, we use accuracy as the only evaluation metric which aligns with the option prediction task. We compute the F_1 score for each question-option pair individually and average across all such pairs. This metric treats every option in every question equally, thus magnifying the impact of options that are rarely selected. For reference, on the Eedi and EdNet datasets, the selection probabilities across options for an average question (from most frequent to least frequent) are 57%, 25%, 11%, 7% and 66%, 20%, 10%, 4%, respectively. For option prediction, a random classifier has an average macro F_1 score and an accuracy score of 0.25 on both of these datasets, while a majority class classifier has an average macro F_1 score (accuracy) of 0.184 (57%) and 0.205 (66%) on the Eedi and EdNet datasets, respectively.

Training and Testing. We perform standard k -fold cross-validation (with $k = 5$) for all methods on both datasets. Under the CF setup, on average 20% of the time steps (for each student) are randomly chosen as the held out test set, 20% of time steps are randomly chosen as the validation set, and the other 60% are chosen as the training set to train all methods. Under the KT setup, all time steps for a randomly chosen 20% of students are used as the test set, and the validation and training sets are constructed similarly.

Network Architectures and Hyper-Parameters. Since the datasets are large, we do minimal hyper-parameter tuning and set most of the values to their default values for all the methods; exploratory experiments found that evaluation results are robust across most parameter values. We set the question, subject, option, response embedding dimension for all methods to $d \in \{32, 64\}$ for the

CF setup and $d \in \{64, 128\}$ for the KT setup. We use the Adam optimizer [22] to train all models with a batch size of 64 students to ensure that an entire batch can fit into the memory of our machine (equipped with one NVIDIA Titan X GPU). For all methods, we set the learning rate to $10^{-4}/10^{-3}$ for the Eedi/EdNet dataset and run all the methods for 200 epochs and perform early stopping based on the loss on the validation set. We set the latent knowledge state (\mathbf{h}_t) dimension to 256/512 for all methods under CF/KT setup. For NCF, we select the user embedding dimension as $d = 256$.

Results and Discussion. Table 1 lists the performance of all OT methods for both the CF and KT setups for both the option prediction and correctness prediction tasks, on both datasets; we report the averages as well as the standard deviations across the five folds. We observe a significant dropoff ($\sim 10\%$) in the accuracy metric on the option prediction task compared to the correctness prediction task, which is as expected since there are four categories to predict (A, B, C, D) instead of two categories (correct/incorrect). As a result of this difference, the correctness prediction task can be seen as a sub-task in the option prediction task by computing the probability a student selects the correct option. The performance of different methods are also quite consistent across all cases.

We observe that recurrent neural network-based methods such as PO-BiDKT perform significantly better than NCF in all cases. This observation suggests that even in a CF setup for model evaluation, methods that take the evolving nature of student knowledge into account are still more effective than popular CF methods that do not account for these temporal dynamics. Overall, we observe that the performance gains on the option prediction task provided by complex model architectures are marginal. This observation suggests that more work needs to be done on the option prediction task to understand the dynamics behind students’ decisions to select a specific incorrect option, which motivates our exploration in Sect. 5.1. In the KT setup, we observe that DKVMN performs best on the Eedi dataset while AKT performs best on the EdNet dataset. This observation suggests that complex neural network architectures such as attention modules are more beneficial when a large amount of training data is available.

In both setups, we observe that the F_1 scores are low for all methods; despite clearly not simply predicting the most frequent option, the performance of these methods leaves significant room for improvement due to class imbalance. Possible approaches to improve prediction accuracy for options that are rarely selected include oversampling them [6]; however, since a student’s responses to different questions are not independent data points, how these methods can be applied to the option tracing task is not immediately clear.

5.1 Clustering Incorrect Options

To qualitatively evaluate our option tracing methods, we attempt to group incorrect options across multiple questions into clusters and examine whether question-option pairs in the same cluster correspond to the same underlying error. To this end, we train a modified version of PO-BiDKT on the Eedi dataset [43]; we learn an embedding module $E_{q,o}(q, o) : q \times o \rightarrow \mathbb{R}^d$ for

Table 2. Incorrect option clustering quality for a subset of questions in the Eedi dataset using errors labeled by a domain expert. 1 in both metrics indicates perfect clustering.

Metric	Adjusted rand index	Fowlkes-Mallows index
Score	0.372	0.455

each question-option pair. Then, we compute the option selection probabilities using the latent knowledge state h_t^i and the question-option pair embeddings as $p(o|q_t^i) = \frac{f(h_t^i)^T E_{q,o}(q_y^i, o)}{\sum_{o'} f(h_t^i)^T E_{q,o}(q_y^i, o')}$. This modification suffers a small drop in predictive performance but encodes information in the question-option pair embeddings for us to cluster them and search for common student errors.

We selected all incorrect options ($31 \times 3 = 63$) in questions on subject 33 (BIDMAS) where question images are released on the smaller subset of the Eedi dataset; see [43] for details. A domain expert manually labeled each option based on which error likely resulted in the student selecting it, resulting in a total of 14 high-level errors (errors that cannot be named are excluded), each corresponding to multiple options across different questions; further splitting them into finer-grained errors results in clusters that are not meaningful. We perform k-means clustering [29] on the learned question-option pair embeddings and compare them to the “ground truth” option clusters provided by the expert.

Due to spatial constraints, we only report quantitative results on clustering quality using two commonly used metrics: The adjusted Rand index [39] and the Fowlkes-Mallows index [14]; the former has a range of $[-1, 1]$ while the latter has a range of $[0, 1]$, with 1 corresponding to perfect clustering.

Table 2 lists these metrics on the learned question-option pair embeddings based on the ground truth expert labeling. Overall, the clustering performance is acceptable but not excellent. We observe that some errors such as “sign error in calculation involving negative numbers” have relatively easy-to-identify corresponding option clusters (5 out of 8 options labeled by the expert as corresponding to that error are put into the same cluster). On the other hand, some options such as 69D and 293C (the left half of Fig. 1) correspond to the same error but are not grouped into the same cluster. One possible explanation is that students may not consistently demonstrate an error, as found in prior research [41]; among students who selected 69D, only 51% selected 293C while 34% of them selected the correct option, 293B. Therefore, further work is required to study whether more robust KT methods and clustering algorithms can identify error clusters more effectively. Nevertheless, our approach produces a starting point to reduce the effort for domain experts to manually label errors and provides them a way to do it under data-driven support.

6 Conclusions and Future Work

Analyzing the exact options students select across multiple choice questions has the potential to uncover their error modes and help teachers to provide targeted

feedback to improve learning outcomes. In this paper, we proposed a set of methods to extend common knowledge tracing methods that analyze only the correctness of students' responses to questions to analyze the exact options they select on multiple choice questions. We validated these methods with quantitative experiments on two large-scale datasets in terms of their ability to predict the options students select on each question and qualitative experiments in terms of clustering incorrect options according to underlying errors. There are many avenues for future work. First, we need to develop methods that are aware of the evolving nature of student errors. One possible approach is to develop methods that can explicitly account for the recurrence of past errors, such as using a neural copy mechanism [18]; these methods may help us track students' progress in correcting their errors. Second, low F_1 scores for the option prediction task suggest that it is much more challenging than the typical correctness prediction task in knowledge tracing literature and thus deserves more attention.

References

1. Adams, D.M., et al.: Using erroneous examples to improve mathematics learning with a web-based tutoring system. *Comput. Hum. Behav.* **36**, 401–411 (2014)
2. Anderson, J.R., Jeffries, R.: Novice LISP errors: undetected losses of information from working memory. *Hum.-Comput. Interact.* **1**(2), 107–131 (1985)
3. Berg, R., Kipf, T.N., Welling, M.: Graph convolutional matrix completion. arXiv preprint [arXiv:1706.02263](https://arxiv.org/abs/1706.02263) (2017)
4. Brown, J.S., Burton, R.R.: Diagnostic models for procedural bugs in basic mathematical skills. *Cogn. Sci.* **2**(2), 155–192 (1978)
5. Cen, H., Koedinger, K., Junker, B.: Learning factors analysis – a general method for cognitive model evaluation and improvement. In: Ikeda, M., Ashley, K.D., Chan, T.-W. (eds.) *ITS 2006*. LNCS, vol. 4053, pp. 164–175. Springer, Heidelberg (2006). <https://doi.org/10.1007/11774303.17>
6. Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: SMOTE: synthetic minority over-sampling technique. *J. Artif. Intell. Res.* **16**, 321–357 (2002)
7. Choffin, B., Popineau, F., Bourda, Y., Vie, J.J.: DAS3H: modeling student learning and forgetting for optimally scheduling distributed practice of skills. In: *Proceedings of the International Conference on Educational Data Mining*, pp. 29–38 (2019)
8. Choi, Y., et al.: EdNet: a large-scale hierarchical dataset in education. In: Bitten-court, I.I., Cukurova, M., Muldner, K., Luckin, R., Millán, E. (eds.) *AIED 2020*. LNCS (LNAI), vol. 12164, pp. 69–73. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-52240-7_13
9. Corbett, A., Anderson, J.: Knowledge tracing: modeling the acquisition of procedural knowledge. *User Model. User-Adapt. Interact.* **4**(4), 253–278 (1994). <https://doi.org/10.1007/BF01099821>
10. Doroudi, S., Aleven, V., Brunskill, E.: Where's the reward? *Int. J. Artif. Intell. Educ.* **29**(4), 568–620 (2019). <https://doi.org/10.1007/s40593-019-00187-x>
11. Erickson, J.A., Botelho, A.F., McAteer, S., Varatharaj, A., Heffernan, N.T.: The automated grading of student open responses in mathematics. In: *Proceedings of the International Conference on Learning Analytics & Knowledge*, pp. 615–624 (2020)

12. Feldman, M.Q., Cho, J.Y., Ong, M., Gulwani, S., Popović, Z., Andersen, E.: Automatic diagnosis of students' misconceptions in K-8 mathematics. In: Proceedings of the CHI Conference on Human Factors in Computing Systems, pp. 1–12 (2018)
13. Feng, J., Zhang, B., Li, Y., Xu, Q.: Bayesian diagnosis tracing: application of procedural misconceptions in knowledge tracing. In: Isotani, S., Millán, E., Ogan, A., Hastings, P., McLaren, B., Luckin, R. (eds.) AIED 2019. LNCS (LNAI), vol. 11626, pp. 84–88. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-23207-8_16
14. Fowlkes, E.B., Mallows, C.L.: A method for comparing two hierarchical clusterings. *J. Am. Stat. Assoc.* **78**(383), 553–569 (1983)
15. Ghosh, A., Heffernan, N., Lan, A.S.: Context-aware attentive knowledge tracing. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 2330–2339 (2020)
16. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge (2016)
17. Graves, A., Mohamed, A.R., Hinton, G.: Speech recognition with deep recurrent neural networks. In: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 6645–6649 (2013)
18. Gu, J., Lu, Z., Li, H., Li, V.O.: Incorporating copying mechanism in sequence-to-sequence learning. arXiv preprint [arXiv:1603.06393](https://arxiv.org/abs/1603.06393) (2016)
19. Gusukuma, L., Bart, A.C., Kafura, D., Ernst, J.: Misconception-driven feedback: results from an experimental study. In: Proceedings of the ACM Conference on International Computing Education Research, pp. 160–168 (2018)
20. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., Chua, T.S.: Neural collaborative filtering. In: Proceedings of the International Conference on World Wide Web, pp. 173–182 (2017)
21. Khajah, M., Huang, Y., González-Brenes, J., Mozer, M., Brusilovsky, P.: Integrating knowledge tracing and item response theory: a tale of two frameworks. In: Proceedings of the International Workshop on Personalization Approaches in Learning Environments, vol. 1181, pp. 7–15 (2014)
22. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: Proceedings of the International Conference on Learning Representations (2015)
23. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint [arXiv:1609.02907](https://arxiv.org/abs/1609.02907) (2016)
24. Koedinger, K.R., Corbett, A., et al.: Cognitive tutors: technology bringing learning sciences to the classroom. In: The Cambridge Handbook of the Learning Sciences, pp. 61–77 (2006)
25. Lan, A.S., Studer, C., Baraniuk, R.G.: Matrix recovery from quantized and corrupted measurements. In: IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 4973–4977, May 2014
26. Lan, A.S., Studer, C., Waters, A.E., Baraniuk, R.G.: Tag-aware ordinal sparse factor analysis for learning and content analytics. In: Proceedings of the 6th International Conference on Educational Data Mining, pp. 90–97, July 2013
27. Lan, A.S., Vats, D., Waters, A.E., Baraniuk, R.G.: Mathematical language processing: automatic grading and feedback for open response mathematical questions. In: Proceedings of the ACM Conference on Learning at Scale, pp. 167–176 (2015)
28. Lindsey, R., Shroyer, J., Pashler, H., Mozer, M.: Improving students' long-term knowledge retention through personalized review. *Psychol. Sci.* **25**(3), 639–647 (2014)
29. Lloyd, S.: Least squares quantization in PCM. *IEEE Trans. Inf. Theory* **28**(2), 129–137 (1982)

30. Ostini, R., Nering, M.L.: Polytomous Item Response Theory Models, No. 144, Sage (2006)
31. Pandey, S., Karypis, G.: A self attentive model for knowledge tracing. In: Proceedings of the International Conference on Educational Data Mining, pp. 384–389, July 2019
32. Pandey, S., Srivastava, J.: RKT: relation-aware self-attention for knowledge tracing. arXiv preprint [arXiv:2008.12736](https://arxiv.org/abs/2008.12736) (2020)
33. Pardos, Z.A., Heffernan, N.T.: Modeling individualization in a Bayesian networks implementation of knowledge tracing. In: De Bra, P., Kobsa, A., Chin, D. (eds.) UMAP 2010. LNCS, vol. 6075, pp. 255–266. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13470-8_24
34. Pavlik Jr., P., Cen, H., Koedinger, K.: Performance factors analysis-a new alternative to knowledge tracing. In: Proceedings of the International Conference on Artificial Intelligence in Education (2009)
35. Piech, C., et al.: Deep knowledge tracing. In: Proceedings of the Conference on Advances in Neural Information Processing Systems, pp. 505–513 (2015)
36. Ritter, S., Anderson, J.R., Koedinger, K.R., Corbett, A.: Cognitive tutor: applied research in mathematics education. *Psychon. Bull. Rev.* **14**(2), 249–255 (2007). <https://doi.org/10.3758/BF03194060>
37. Selent, D.A.: Creating systems and applying large-scale methods to improve student remediation in online tutoring systems in real-time and at scale. Ph.D. thesis, Worcester Polytechnic Institute (2017)
38. Smith III, J.P., DiSessa, A.A., Roschelle, J.: Misconceptions reconceived: a constructivist analysis of knowledge in transition. *J. Learn. Sci.* **3**(2), 115–163 (1994)
39. Steinley, D.: Properties of the Hubert-arable adjusted rand index. *Psychol. Methods* **9**(3), 386 (2004)
40. Thissen, D., Steinberg, L.: A taxonomy of item response models. *Psychometrika* **51**(4), 567–577 (1986)
41. VanLehn, K.: Bugs are not enough: empirical studies of bugs, impasses and repairs in procedural skills. *J. Math. Behav.* (1982)
42. Wang, F., et al.: Neural cognitive diagnosis for intelligent education systems. *Pro. AAAI Conf. Artif. Intell.* **34**, 6153–6161 (2020)
43. Wang, Z., et al.: Diagnostic questions: the NeurIPS 2020 education challenge. arXiv preprint [arXiv:2007.12061](https://arxiv.org/abs/2007.12061) (2020)
44. Woolf, B.P.: Building Intelligent Interactive Tutors: Student-Centered Strategies for Revolutionizing E-learning. Morgan Kaufmann (2010)
45. Yang, Y., et al.: GIKT: a graph-based interaction model for knowledge tracing. In: Hutter, F., Kersting, K., Lijffijt, J., Valera, I. (eds.) ECML PKDD 2020. LNCS (LNAI), vol. 12457, pp. 299–315. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-67658-2_18
46. Yudelson, M.V., Koedinger, K.R., Gordon, G.J.: Individualized Bayesian knowledge tracing models. In: Lane, H.C., Yacef, K., Mostow, J., Pavlik, P. (eds.) AIED 2013. LNCS (LNAI), vol. 7926, pp. 171–180. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39112-5_18
47. Zhang, J., Shi, X., King, I., Yeung, D.Y.: Dynamic key-value memory networks for knowledge tracing. In: Proceedings of the International Conference on World Wide Web, pp. 765–774, April 2017