

# Assessing the Performance of Online Students - New Data, New Approaches, Improved Accuracy

Robin Schmucker  
Carnegie Mellon University  
rschmuck@cs.cmu.edu

Jingbo Wang  
Carnegie Mellon University  
jingbow@andrew.cmu.edu

Shijia Hu  
Carnegie Mellon University  
shijiah@andrew.cmu.edu

Tom M. Mitchell  
Carnegie Mellon University  
tom.mitchell@cs.cmu.edu

---

We consider the problem of assessing the changing performance levels of individual students as they go through online courses. This student performance modeling problem is a critical step for building adaptive online teaching systems. Specifically, we conduct a study of how to utilize various types and large amounts of log data from earlier students to train accurate machine learning models that predict the performance of future students. This study is the first to use four very large sets of student data made available recently from four distinct intelligent tutoring systems.

Our results include a new machine learning approach that defines a new state of the art for logistic regression based student performance modeling, improving over earlier methods in several ways: First, we achieve improved accuracy of student modeling by introducing new features that can be easily computed from conventional question-response logs (e.g., features such as the pattern in the student's most recent answers). Second, we take advantage of features of the student history that go beyond question-response pairs (e.g., features such as which video segments the student watched, or skipped) as well as background information about prerequisite structure in the curriculum. Third, we train multiple specialized student performance models for different aspects of the curriculum (e.g., specializing in early versus later segments of the student history), then combine these specialized models to create a group prediction of the student performance. Taken together, these innovations yield an average AUC score across these four datasets of 0.808 compared to the previous best logistic regression approach score of 0.767, and also outperforming state-of-the-art deep neural net approaches. Importantly, we observe consistent improvements from each of our three methodological innovations, in each diverse dataset, suggesting that our methods are of general utility and likely to produce improvements for other online tutoring systems as well.

**Keywords:** performance modeling, knowledge tracing, logistic regression, deep learning, features

---

## 1 INTRODUCTION

Intelligent online tutoring systems (ITS's) are now used by millions of students worldwide, enabling access to quality teaching materials and to personally customized instruction. These

systems depend critically on their ability to track the evolving ability level of the student, in order to deliver the most effective instructional material at each point in time. Because the problem of assessing the student's evolving ability to solve different questions (also referred to as student performance modeling) is so central to successfully customizing instruction to the individual student, it has received significant attention in recent years.

The state of the art that has emerged for this student performance modeling problem involves applying machine learning algorithms to historical student log data. The result produced by the machine learning algorithm is a model, or computer program, that outputs the estimated likelihood of correct response of any future student for any particular question at any point in the lesson, given the sequence of steps they have taken up to this point in the lesson (e.g., [Corbett and Anderson 1994](#), [Pavlik Jr et al. 2009](#), [Piech et al. 2015](#), [Pandey and Karypis 2019](#), [Gervet et al. 2020](#), [Shin et al. 2021](#)). These systems typically represent the student knowledge state as a list of probabilities that the student will correctly answer a particular list of questions that cover the key concepts (also known as "knowledge components" (KC's)) in the curriculum. Most current approaches estimate the student's state by considering only the log of questions asked and the student's answers, though recent datasets provide considerably more information such as the length of time taken by the student to provide answers, specific videos the student watched and whether they watched the entire video, and what hints they were given as they worked through specific practice problems.

This paper seeks to answer the question of which machine learning approach produces the most accurate estimates of students' ability to solve different questions. To answer this question, we perform an empirical study using data from over 750,000 students taking a variety of courses, to study several aspects of the question including (1) which types of machine learning algorithms work best? (2) which features of a student's previous and current interactions with the ITS are most useful for predicting their current ability to solve a certain question? (3) how valuable is background information about curriculum prerequisites for improving accuracy? and (4) can accuracy be improved by training specialized models for different portions of the curriculum? We measure the quality of alternative approaches by how accurately they predict which future questions the student answers correctly.

More specifically, we present here the first comparative analysis of recent state-of-the-art algorithms for student performance modeling across four very large student log datasets that have recently become available, which are each approximately 10 times larger than earlier publicly available datasets, which cover a variety of courses in elementary mathematics, as well as teaching English as a second language, and which range across different teaching objectives such as initial assessment of student knowledge state, test preparation, and extra-curricular tutoring complementing K-12 schooling. We show that accuracy of student performance modeling can be improved beyond the current state of the art through a combination of techniques including incorporating new features from student logs (e.g., time spent on previously answered questions), incorporating background information about prerequisite/post-requisite topics in the curriculum, and training multiple specialized models for different parts of the student experience (e.g., training distinct models to assess new students during the first 10 steps of their lesson, versus students taking the post-lesson quiz). The fact that we see consistent improvements in accuracy across all four datasets suggests the lessons gained from our experiments are fairly general, and not tied to a specific type of course or specific tutoring system.

To summarize, the key contributions of this paper include:

- *Cross-ITS study on modern datasets.* We present the first comparative analysis of state-of-the-art approaches to student performance modeling across four recently published, large and diverse student log datasets taken from four distinct intelligent tutoring systems (ITS's), resulting in the largest empirical study to date of student performance modeling. These four systems teach various topics in elementary mathematics or English as a second language, and the combined data covers approximately 200,000,000 observed actions taken by approximately 750,000 students. Three of these datasets have been made publicly available over the past few years.
- *Improved student performance modeling by incorporating prerequisite and hierarchical structure across knowledge components.* All four of the datasets we consider provide a vocabulary of knowledge components (KCs) taught by the system. Two datasets provide meta-information about which KCs are prerequisites for which others (e.g., *Add and subtract negative numbers* is a prerequisite for *Multiply and divide negative numbers*), and one provides a hierarchical structure (e.g., the KC *Add and subtract vectors* falls hierarchically under *Basic vectors*). We found that incorporating the prerequisite and post-requisite information into the model resulted in substantial improvements in accuracy predicting which future questions the student would answer correctly (e.g., by including features such as the counts of correctly answered questions related to prerequisite and post-requisite topics). Features derived from the hierarchical structure led to improvements as well, but to a lesser degree than the ones extracted from the prerequisite structures.
- *Improved student performance modeling by incorporating log data features beyond which questions were answered correctly.* Although earlier work has focused on student performance models that consider only the sequence of questions asked and which were answered correctly, modern log data includes much more information. We found that incorporating this information yields accuracy improvements over the current state of the art. For example, we found that features such as the length of time the student took to answer previous questions, the number of videos watched on the KC, and whether the student is currently answering a question in a pre-test, post-test, or a practice session were all useful features for predicting whether a student would correctly answer the next question.
- *Improved student performance modeling by training multiple models for specialized contexts, then combining them.* We introduce a new approach of training multiple distinct student assessment models for distinct learning contexts. For example, training a distinct model for the "cold start" problem of assessing students who have just begun the course and have little log data at this point, yields significant improvements in the accuracy of student performance predictions. Furthermore, combining predictions of multiple specialized models (e.g., one trained for the current study module type, and one trained for students that have answered at least  $n$  questions in the course) leads to even further improvements.
- *The above improvements can be combined.* The above results show improvements in student performance modeling due to multiple innovations. By combining these we achieve overall improvements over state-of-the-art logistic regression methods that reduce AUC error by 17.5% on average over these four datasets, as summarized in Table 1.

Table 1: Improvements to state of the art (SOTA) in student performance modeling, due to innovations introduced in this paper. The previous logistic regression state-of-the-art approach is Best-LR. Performance across each of the four diverse datasets improves with each of our three suggested extensions to the Best-LR algorithm. Best-LR+ extends Best-LR by adding new features calculated from the question-response (Q-R) data available in most student logs. AugmentedLR further adds a variety of novel features that go beyond question-response data (e.g., which videos the student watched or skipped). Combined-AugmentedLR further extends the approach by training multiple logistic regression models (Multimodel) on different subsets of the training data (e.g., based on how far into the course the student is currently). Together, these extensions to the previous state of the art produce substantial improvements across each of these four datasets, improving the average AUC score from 0.767 to 0.808 on average across the four datasets – a reduction of 17.5% in average AUC error (i.e., in the difference between the observed AUC and the ideal perfect AUC of 1.0).

	ElemMath2021		EdNet KT3		Eedi		Junyi15		Average	
	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC
Prev. SOTA: Best-LR	0.7569	0.7844	0.7069	0.7294	0.7343	0.7901	0.8425	0.7620	0.7602	0.7665
Add Q-R features: Best-LR+	0.7623	0.7935	0.7169	0.7465	0.7455	0.8040	0.8505	0.7912	0.7688	0.7838
Add novel features: AugmLR	0.7659	0.7987	0.7189	0.7500	0.7496	0.8096	0.8635	0.8603	0.7745	0.8047
Multimodel: Comb-AugmLR	<b>0.7676</b>	<b>0.8016</b>	<b>0.7211</b>	<b>0.7548</b>	<b>0.7504</b>	<b>0.8111</b>	<b>0.8646</b>	<b>0.8634</b>	<b>0.7759</b>	<b>0.8077</b>
percent error reduction	4.40%	7.98%	4.84%	9.39%	6.06%	10.00%	14.03%	42.61%	7.33%	17.50%

## 2 RELATED WORK

Student performance modeling techniques estimate a student’s likelihood to solve different problems based on their interactions with the ITS. These performance models and the estimates of student proficiency they produce are a key component of current ITS’s which allow the tutoring system to adapt to each student’s personal ability level at each point in the curriculum. In the literature, performance modeling is also sometimes referred to as knowledge tracing, proficiency modeling, or student assessment. There are three main categories of performance modeling techniques: (i) Markov process based probabilistic modeling, (ii) logistic regression, and (iii) deep learning based approaches.

Markov process based techniques, such as Bayesian Knowledge Tracing (BKT) (Corbett and Anderson, 1994) and its various extensions (Baker et al. 2008, Pardos and Heffernan 2010, Pardos and Heffernan 2011, Qiu et al. 2011, Yudelson et al. 2013, Sao Pedro et al. 2013, Khajah et al. 2016, Käser et al. 2017) have a long history in the educational data mining (EDM) community. Most approaches in this family determine a student proficiency by performing probabilistic inference using a two state Hidden Markov Model containing one state representing that the student has *mastered* a particular concept, and one state representing *non-mastery*. A recent study comparing various performance modeling algorithms across nine real-world datasets (Gervet et al., 2020) found that when applied to large-scale datasets, BKT and its extension BKT+ (Khajah et al., 2016) are very slow to train and their predictive performance is not competitive with more recent logistic regression and deep learning based approaches. The Python package pyBKT was released and promises faster training times for Bayesian Knowledge Tracing models (Badrinath et al., 2021).

Logistic regression models take as input a vector of manually specified features calculated from a student’s interaction history, then output a predicted probability that this student has mastered a particular concept or KC (often implemented as the probability that they will correctly

answer a specified question). Common approaches include IRT ([van der Linden and Hambleton, 2013](#)), LFA ([Cen et al., 2006](#)), PFA ([Pavlik Jr et al., 2009](#)), DASH ([Lindsey et al. 2014](#), [González-Brenes et al. 2014](#), [Mozer and Lindsey 2016](#)) and its extension DAS3H ([Choffin et al., 2019](#)) as well as Best-LR ([Gervet et al., 2020](#)). While there exists a variety of logistic regression models, they mainly rely on two types of features: (i) One-hot encodings<sup>1</sup> of question and KC identifiers; and (ii) Count features capturing the student’s number of prior attempts to answer questions, and the number of correct and incorrect responses. R-PFA ([Galyardt and Goldin, 2015](#)) augments PFA with features that represent the recency-weighted count of prior incorrect responses and the recency-weighted proportion of correct responses. PPE ([Walsh et al., 2018](#)) considers the timing of individual practice sessions to describe spacing effects impacting memorization in the context of learning word pairs via power functions. DAS3H incorporates a temporal aspect into its predictions by computing count features for different time windows. LKT ([Pavlik et al., 2021](#)) is a flexible framework for logistic regression based student performance modeling which offers a variety of features based on question-answering behavior. Among others it offers decay functions to capture recency effects as well as features based on the ratio of correct and incorrect responses. Section 4 discusses multiple regression models in more detail and proposes alternative features which are able to incorporate rich information from various types of log data.

Deep learning based models, like logistic regression models, take as input the student log data, and output a predicted probability that the student will answer a specific question correctly. However, unlike logistic regression, deep learning models have the ability to automatically define useful features computable from the sequence of log data, without relying on human feature engineering. A wide range of neural architectures have been proposed for student performance modeling. DKT ([Piech et al., 2015](#)) is an early work that uses Long Short Term Memory (LSTM) networks ([Hochreiter and Schmidhuber, 1997](#)) processing student interactions step-by-step. DKVNM ([Zhang et al., 2017](#)) is a memory augmented architecture that can capture multi-KC dependencies. CKT ([Shen et al., 2020](#)) uses a convolutional neural network ([LeCun et al., 1999](#)) to model individualized learning rates. Inspired by recent advances in the natural language processing (NLP) community, multiple transformer ([Vaswani et al., 2017](#)) based approaches have been proposed (SAKT, ([Pandey and Karypis, 2019](#)); AKT, ([Ghosh et al., 2020](#)); SAINT, ([Choi et al., 2020](#)); SAINT+, ([Shin et al., 2021](#))). Graph-based modeling approaches infer the likelihood with which a certain problem is answered correctly based on the structure induced by question-KC relations (GKT, ([Nakagawa et al., 2019](#)); HGKT, ([Tong et al., 2020](#)); GIKT, ([Yang et al., 2021](#))). Unlike BKT techniques and certain logistic regression based approaches such as IRT, deep learning based models often do not provide interpretable parameters or predictions which allow users to quantify a learner’s knowledge related to a particular KC. To mitigate this shortcoming recent works have proposed more interpretable network architectures ([Yeung 2019](#), [Tsutsumi et al. 2021](#)) and techniques to derive knowledge estimates directly from performance predictions ([Scruggs et al., 2020](#)). For an in detail survey on the recent student performance modeling approaches we refer to [Liu et al. 2021](#).

Most student performance modeling approaches focus exclusively on question-answering behavior. A student’s sequence of past interactions with the tutoring system is modelled as  $\mathbf{x}_{1:t} = (x_1, \dots, x_t)$ . The  $t^{\text{th}}$  response is represented by a tuple  $x_t = (q_t, a_t)$ , where  $q_t$  is the question (item) identifier and  $a_t \in \{0, 1\}$  is binary response correctness. While this formalism

---

<sup>1</sup>A one-hot encoding of the question ID, for example, is a vector whose length is equal to the number of possible questions. The vector contains  $n - 1$  zeros, and a single 1 to indicate which question ID is being encoded.



has yielded many effective performance modeling techniques, it is limiting in that it assumes the student log data contains only a single type of user interaction: answering questions posed by the ITS. Many aspects of student behavior such as interactions with learning materials (videos, instructional text), hint usage, and information about the current learning context can provide a useful signal, but fall outside the scope of this formalism.

More recently, the EDM community has started exploring alternative types of log data, in an attempt to improve the accuracy of student performance modeling. [Zhang et al. 2017](#) augment DKT with information on response time, attempt number and type of first interaction with the ITS. Later, [Yang and Cheung 2018](#) enhanced DKT predictions by incorporating more than 10 additional features related to learning context, interaction times and hint usage provided by the ASSISTment 2009 ([Feng et al., 2009](#)) and Junyi15 ([Chang et al., 2015](#)) datasets. While that work employs most information contained in the Junyi15 dataset, it fails to utilize the prerequisite structure among topics in the curriculum, and does not evaluate the potential benefit of those features for logistic regression models. EKT ([Liu et al., 2019](#)) uses the question text to learn exercise embeddings which are then used for downstream performance predictions. [Eglington and Pavlik 2019](#) use student response times and correctness rates to cluster the user population into multiple groups each representing a different student phenotype. They then incorporate a set of cluster specific model parameters into a modified PFA model to capture variations between the individual groups leading to improved performance predictions. MVKM ([Zhao et al., 2020](#)) uses a multi-view tensor factorization to model knowledge acquisition from different types of learning materials (e.g., quizzes, videos, ...). A recent line of research identified a *Doer* effect which is associated with the finding that interactive problem solving is more indicative for learning outcomes than more passive study activities such as reading and watching lecture videos ([Koedinger et al. 2015](#), [Koedinger et al. 2016](#), [Koedinger et al. 2018](#), [Van Campenhout et al. 2021](#)). SAINT+ ([Shin et al., 2021](#)) and MUSE ([C. Zhang et al., 2021](#)) augment transformer models with interaction time features to capture short-term memorization and forgetting. Closest to the spirit of this manuscript is an early work by [Feng et al. 2009](#). That work integrates features related to accuracy, response time, attempt-usage and help-seeking behavior into a logistic regression model to enhance exam score predictions. However, their exam score prediction problem is inherently different from our problem of continuous performance modeling because it focuses only on the final learning outcome and does not capture user proficiency on the question and KC level. Unlike our work, [Feng et al. 2009](#) do not evaluate features related to lecture video and reading material consumption, prerequisite structure and learning context. Their study is also limited to two small-scale datasets (< 1000 students) collected by the ASSISTment system.

This paper offers the first systematic study of a wide variety of features extracted from alternative types of log data. It analyzes four recent large-scale datasets which capture the learning process at different levels of granularity and focus on various aspects of student behavior. Our study identifies a set of features which can improve the accuracy of student performance modeling techniques and we give recommendations regarding which types of user interactions should be captured in log data of future tutoring systems. Further, our feature evaluation led us to novel logistic regression models achieving a new state-of-the-art performance on all four datasets.

Table 2: Summary of datasets. Here *KC* refers to a knowledge component in the curriculum of the respective system, *average correctness* refers to the fraction of questions that were correctly answered across all students, *question ID* refers to whether each distinct question (item) had an ID, and *KC ID* indicates that each item is linked to at least one KC, *platform* indicates if the system logs how students access materials (e.g., mobile app or web browser), *social support* indicates if there is information about a student’s socioeconomic status, *question bundle* indicates if the system groups multiple items into sets which are asked together, *elapsed/lag time* indicates whether the data includes detailed timing information allowing the calculation of how long the student took to respond to each question, and the time between presentations of successive questions. *Question difficulty* indicates whether the dataset includes ITS-defined difficulties for each question. *Videos, reading, and hints* indicate whether the dataset provides information about which videos and explanations the student watched or read, and which hints were delivered to them as they worked on practice questions.

	ElemMath2021	EdNet KT3	Eedi	Junyi15
# of students	125,246	297,915	118,971	247,606
# of unique questions	59,892	13,169	27,613	835
# of KCs	4,191	293	388	41
# of logged actions	62,570,009	89,270,654	19,834,813	25,925,992
# of student responses	23,447,961	17,954,718	19,834,813	25,925,992
average correctness	68.52%	66.19%	64.30%	82.99%
subject	Mathematics	English	Mathematics	Mathematics
timestamp	✓	✓	✓	✓
question ID	✓	✓	✓	✓
KC ID	✓	✓	✓	✓
age/gender			✓	
social support			✓	
platform	✓	✓		
teacher/school	✓		✓	
study module	✓	✓	✓	✓
pre-requisite graph	✓			✓
KC hierarchy			✓	
question bundle		✓	✓	
elapsed/lag time	✓	✓		✓
question difficulty	✓			
videos	✓	✓		
reading	✓	✓		
hints				✓

### 3 DATASETS

In recent years multiple education companies released large-scale ITS datasets to promote research on novel educational data mining techniques. Compared to earlier sizable datasets such as Bridge to Algebra 2006 (Stamper et al., 2010) and ASSISTment 2012 (Feng et al., 2009), these new datasets capture an order of magnitude more student responses making

them attractive for data intensive modeling approaches. Table 2 provides an overview of the four recent large-scale student log datasets, from different ITS systems, we use in this paper. Taken together, these datasets capture over 197 million lines of log data from over 789,000 students, including correct/incorrect responses to over 87 million questions. While each ITS collects the conventional timestamp for each question (item) presented to the student, the correctness of their answer, and knowledge component (KC) attributes associated with the question, they also contain additional interaction data. Going beyond pure question-solving activities multiple datasets provide information about how and when students utilize reading materials, lecture videos and hints. In addition there are various types of meta-information about the individual students (e.g., age, gender, ...) as well as the current learning context (e.g., school, topic number, ...). All four tutoring systems exhibit a modular structure in which learning activities are assigned to distinct categories (e.g., pre-test, effective learning, review, ...). Further, three of the datasets provide meta-information indicating which questions, videos, etc. are related to which KCs.

Figure 2 visualizes the distribution of the number of responses per student for each dataset (i.e., the number of answered questions per user). All datasets follow a power-law distribution. The EdNet<sub>KT3</sub> and Junyi15 datasets contain many users with less than 50 completed questions and only a small proportion of users answer more than 100. The Eedi dataset filters out students and questions with less than 50 responses. The ElemMath2021 dataset exhibits the largest median response number, and many users answer several hundreds of questions.

Another interesting ITS property is the degree to which all students progress through the study material in a fixed sequence, versus how different and adaptive is the presented sequence across different students. We can get insight into this for each of our datasets by asking how predictable the next question item (or the next KC) is from the previous one, across all students in the dataset. We determine predictability by evaluating the accuracy of a simple model which takes as input the ID of the current question or KC and outputs the ID of the next question or KC which is most likely based on the empirical successor distribution as captured by the log data for the input ID. Figure 1 shows how accurately the current question/KC predicts the following question/KC. The Junyi15 data exhibits very low variability in its KC sequencing and also tends to present questions in the same order across all students. Eedi logs show more variability in the KC sequence, but the question order is still rather predictable. ElemMath2021 exhibits moderate KC sequence variations and a highly variable question order. EdNet<sub>KT3</sub> exhibits the most variation in question and KC order.

Some of the differences between the individual datasets can be traced back to the distinct structures and objectives of the underlying tutoring systems. They teach different subjects, assist students in different ways (K-12 tutoring, standardized test preparation, knowledge diagnosis), and provide students with varying levels of autonomy. We next discuss the individual datasets and corresponding ITS in more detail.

**EdNet** (Choi et al., 2020): EdNet was introduced as a large-scale benchmarking dataset for student performance modeling algorithms. The data was collected over two years by Riiid’s Santa tutoring system which prepares students in South Korea for the *Test of English for International Communication* (TOEIC<sup>®</sup>) Listening & Reading. Each test is split into seven distinct parts, four assessing listening and three assessing reading proficiency. Following this structure Santa categorizes its questions into seven parts (additional fine-grained KCs are provided). Santa is a multi-platform system available on Android, iOS and the web and users have autonomy regarding which test parts they want to focus on. There exist four versions of this dataset (KT1, ..., KT4) capturing student behavior at increasing levels of detail ranging from pure question answering






Figure 1: Accuracy when predicting the next question or KC based on the current one. Lower accuracy reflects greater sequence variability across students.

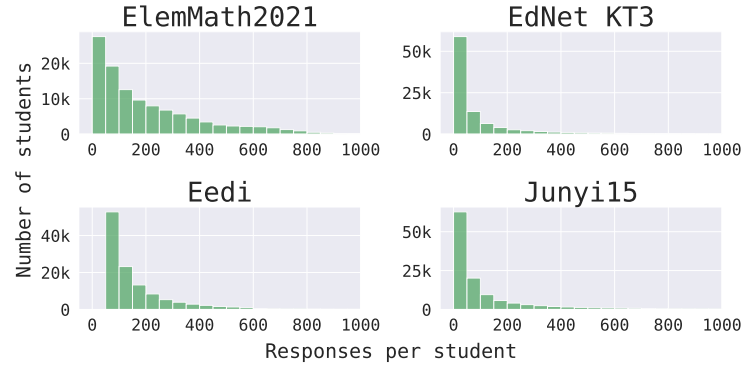


Figure 2: Distribution of the number of responses per student. All four datasets follow a power-law distribution. The EdNet KT3 and Junyi15 dataset both contain many users with few responses. The users of the Eedi and ElemMath2021 tutoring systems have more responses on average.

activity to comprehensive UI interactions. In this paper we analyze the EdNet KT3 dataset which contains logs of 297,915 students and provides access to reading and video consumption behavior. We omit the use of the KT4 dataset which augments KT3 with purchasing behavior.

**Junyi15** (Chang et al., 2015): The Junyi Academy Foundation is a philanthropic organization located in Taiwan. It runs the Junyi Academy online learning platform which offers various educational resources designed for K-12 students. In 2015 the foundation released log data from their mathematics curriculum capturing activities of 247,606 students collected over 2 years. Users of Junyi Academy can choose from a large variety of topics, are able to submit multiple answers to a single question, and can request hints. The system registers an answer as correct if no hints are used and the correct answer is submitted on the first attempt. The dataset provides a prerequisite graph which captures semantic dependencies between the individual questions. In addition to a single KC each question is annotated with an area identifier (e.g., algebra, geometry, ...). In 2020 Junyi Academy shared a newer mathematics dataset on Kaggle (Pojen et al., 2020). Unfortunately, in that dataset each timestamp is rounded to the closest quarter hour which prevents exact reconstruction of the response sequence making it difficult to evaluate student performance modeling algorithms. Because of this we perform our analysis using the Junyi15 dataset.

**Eedi** (Wang et al., 2020): Eedi is an UK based online education company which offers a knowledge assessment and misconception diagnosis service. Unlike other ITS which are designed to teach new skills, the Eedi system confronts each student with a series of *diagnostic questions* – i.e., multiple choice questions in which each incorrect answer is indicative of a common misconception. This process results in a report that helps school teachers to adapt to student specific needs. In 2020 Eedi released a dataset for the NeurIPS Education Challenge (Wang et al., 2021). It contains mathematics question logs of 118,971 students (primary to high school) and was collected over a 2-year period. Student age and gender as well as information on whether a student qualifies for England’s pupil premium grant (a social support program for disadvantaged students) is provided. In contrast to the Junyi15 and ElemMath2021 datasets which have a prerequisite graph, the Eedi dataset organizes its KCs via a 4-level topic ontology tree. For

example, the KC *Add and Subtract Vectors* falls under the umbrella of *Basic Vectors* which itself is assigned to *Geometry and Measure* which is connected to the subject *Mathematics*. While analyzing this dataset we noticed that the timestamp information is rounded to the closest minute which prevents exact reconstruction of the interaction sequences. Upon request, the authors provided us with an updated version of the dataset that allows exact recovery of the interaction sequence.

**Squirrel Ai ElemMath2021:** Squirrel Ai Learning (SQ-Ai) is a K-12 education company located in China which offers individualized after-school tutoring services. SQ-Ai provides their ITS as mobile and Web applications, but also deploys it in over 3000 physical tutoring centers. The centers provide a unique setting in which students can study under the supervision of human teachers and can ask for additional advice and support which augments the ITS's capabilities. Students also have the social experience of working alongside their peers. In this paper we introduce and analyze the Squirrel Ai ElemMath2021 dataset. It provides 3 months of behavioral data of 125,246 K-12 students completing various mathematics courses and captures observational data at fine granularity. ElemMath2021 gives insight into reading material and lecture video consumption. It also provides meta-information on which learning center and teacher a student is associated with. Each question has a manually assigned difficulty rating ranging from 10 to 90. A prerequisite graph captures dependencies between the individual KCs. Most student learning sessions have a duration of about one hour and the ITS selects a session topic. Each learning process is assigned one of six categories and learning sessions usually follow a pre-test, learning, post-test structure. This makes this dataset particularly interesting because it allows to quantify the learning success of each individual session as the difference between pre-test and post-test performance.

We conclude this section with a few summarizing remarks. Recent years have yielded large-scale educational datasets which have the potential to fuel future EDM research. The individual datasets exhibit large heterogeneity with regards to the structure and objectives of the underlying tutoring systems as well as captured aspects of the learning process. Motivated by these observations we perform the first comparative evaluation of various student performance modeling techniques across these four large-scale datasets. Further, we use the rich log data to evaluate potential benefits of a variety of alternative features to provide recommendations on which types of observational data is informative for future performance modeling techniques. Our feature evaluation leads us to multiple novel logistic regression models achieving state-of-the-art performance. Given the size of available training data and the structure of the learning processes we also address the question if it is beneficial to train a set of different assessment modules which are specialized on different parts of the ITS.

## 4 APPROACH

In this paper we examine alternative student performance modeling algorithms using alternative sets of features, across four recent large-scale ITS datasets. Our goals are (1) to discover the features of rich student log data that lead to the most accurate performance models of students, (2) to discover the degree to which the most useful features are task-independent versus dependent on the particular tutoring system, and (3) to discover which types of machine learning algorithms produce the most accurate student models using this log data. We start this Section with a formal definition of the student performance modeling problem. We then discuss prior work on logistic regression based modeling approaches and analyze which types of features they employ. From

there, we introduce a set of alternative features leveraging alternative types of log data to offer a foundation for novel student performance prediction algorithms. We conclude this Section with the proposal of two additional features which capture long-term and short-term student performance and only rely on response correctness. Appendix A provides precise definitions of all features used in our experiments, including implementations details and ITS specific considerations.

#### 4.1 THE STUDENT PERFORMANCE MODELING PROBLEM

Student performance modeling is a supervised sequence learning task which traces a student’s likelihood to solve different problems over time. Reliable performance predictions are crucial to enabling ITSs to provide effective individualized feedback to users. More formally, we denote a student’s sequence of past interactions with the tutoring system as  $\mathbf{x}_{1:t} = (x_1, \dots, x_t)$ . The  $t^{\text{th}}$  interaction with the system is represented by the tuple  $x_t = (I_t, c_t)$ , where  $I_t$  indicates the interaction type and  $c_t$  is a dataset dependent aggregation of information related to the interaction. In this paper we consider interaction types connected to question answering, video and reading material consumption as well as hint usage. Examples of attributes contained in  $c_t$  are timestamp, learning material identifiers, information about the current learning context and student specific features. Question answering is the most basic interaction type which is monitored by all ITSs. If the  $t^{\text{th}}$  interaction is a question response,  $c_t$  provides the question (item) identifier  $q_{t+1}$  and binary response correctness  $a_t \in \{0, 1\}$ . Given a user’s history of past interaction with the ITS, the student performance problem is to predict  $p(a_{t+1} = 1 | q_{t+1}, \mathbf{x}_{1:t})$  – the probability that the student’s response will be correct if they are next asked question  $q_{t+1}$ , given their history  $\mathbf{x}_{1:t}$ . In addition to interaction logs, all four datasets provide a knowledge component model which associates each question  $q_t$  with a set  $KC(q_t)$  containing one or more knowledge components (KCs). Each KC represent a concrete skill which can be targeted by questions and other learning materials. User interactions are discrete and observed at irregular time intervals. To capture short-term memorization and forgetting it is necessary to utilize additional temporal features. We denote the dependence of variables on the individual student with the subscript  $s$ .

#### 4.2 LOGISTIC REGRESSION FOR STUDENT PERFORMANCE MODELING

Logistic regression models enjoy great popularity in the EDM community. At their core each trained regression model takes as input a real-valued feature vector  $(\phi_1, \dots, \phi_d)$  that describes the student  $s$  and their log data up to this point in the course, along with a question ID. Note different logistic regression approaches can summarize the student and their history in terms of different features. Each approach calculates its features using its own feature calculation function,  $\Phi$  (i.e.,  $(\phi_1, \dots, \phi_d) = \Phi(q_{s,t+1}, \mathbf{x}_{s,1:t})$ ). The trained logistic regression model then uses this feature vector as input, and outputs the probability that the student will correctly answer question  $q_{s,t+1}$  if they are asked it at this point in time. The full logistic regression model is therefore of the form

$$p(a_{s,t+1} = 1 | q_{s,t+1}, \mathbf{x}_{s,1:t}) = \sigma(\mathbf{w}^\top \Phi(q_{s,t+1}, \mathbf{x}_{s,1:t})). \quad (1)$$

Here  $\mathbf{w} \in \mathbb{R}^d$  represents the vector of learned regression weights and  $\sigma(x) = 1/(1 + e^{-x})$  is the sigmoid function which outputs a value between 0 and 1, which is interpreted as the probability that student  $s$  will answer question  $q_{t+1}$  correctly. A suitable set of weights can be determined by maximizing the likelihood function on the training set. The corresponding maximization problem is usually solved using gradient based optimization algorithms.

Over the years various logistic regression models have been proposed, each employing a distinct feature mapping  $\Phi$  to extract a suitable feature set. In this paper we consider Item Response Theory (IRT; [van der Linden and Hambleton 2013](#)), in particular a one-parameter version of IRT known as Rasch model ([Rasch, 1960](#)), Performance Factor Analysis (PFA; [Pavlik Jr et al. 2009](#)), Recent-Performance Factors Analysis (R-PFA; [Galyardt and Goldin 2015](#)), Predictive Performance Equation (PPE; [Walsh et al. 2018](#)), DAS3H ([Choffin et al., 2019](#)) and Best-LR ([Gervet et al., 2020](#)). We now discuss the individual models focusing on how the available interaction data is utilized for their predictions. First, IRT is the simplest regression model. It employs a parameter  $\alpha_s$  which represents the ability of student  $s$  as well as a separate difficulty parameter  $\delta_q$  for each question (item)  $q$ . The IRT prediction is defined as

$$p_{\text{IRT}}(a_{s,t+1} = 1 \mid q_{s,t+1}, \mathbf{x}_{s,1:t}) = \sigma(\alpha_s - \delta_{q_{s,t+1}}). \quad (2)$$

Unlike IRT, PFA extracts features based on a student’s history of past interactions. It computes the number of correct ( $c_{s,k}$ ) and incorrect responses ( $f_{s,k}$ ) prior to the current attempt and introduces a difficulty parameter  $\beta_k$  for each individual KC  $k$ . The PFA prediction is defined as

$$p_{\text{PFA}}(a_{s,t+1} = 1 \mid q_{s,t+1}, \mathbf{x}_{s,1:t}) = \sigma\left(\sum_{k \in \text{KC}(q_{s,t+1})} \beta_k + \gamma_k c_{s,k} + \rho_k f_{s,k}\right). \quad (3)$$

R-PFA is motivated by the idea that more recently observed student responses are more indicative for future performance than older ones. R-PFA builds on PFA by introducing two features which for each KC  $k$  look at all interactions of student  $s$  with  $k$  up to time  $t$  and computes: (i) A recency-weighted count of previous failures  $F_{s,k,t}$  using exponential decay. (ii) A recency-weighted proportion of past successes  $R_{s,k,t}$  using normalized exponential decay. The degree of decay is controlled by the hyperparameters  $d_F$  and  $d_R \in [0, 1]$ . To allow the computation of  $R_{s,k,t}$  when a student visits a KC for the first time, R-PFA appends their interaction history with  $k$  with  $g = 3$  incorrect “ghost attempt”. We denote the total number of responses of student  $s$  related to KC  $k$  as  $a_{s,k}$  and use a correctness indicator  $a_{s,k,i}$  which is 1 when  $s$ ’s  $i$ -th attempt on KC  $k$  was correct and 0 otherwise. The R-PFA prediction is defined as

$$p_{\text{R-PFA}}(a_{s,t+1} = 1 \mid q_{s,t+1}, \mathbf{x}_{s,1:t}) = \sigma\left(\sum_{k \in \text{KC}(q_{s,t+1})} \beta_k + \gamma_k F_{s,k,t} + \rho_k R_{s,k,t}\right) \quad (4)$$

$$F_{s,k,t} = \sum_{i=1}^{a_{s,k}} d_F^{(a_{s,k}+1)-i} (1 - a_{s,k,i}), \quad R_{s,k,t} = \sum_{i=(1-g)}^{a_{s,k}} \frac{d_R^{a_{s,k}-i}}{\sum_{j=(1-g)}^{a_{s,k}} d_R^{a_{s,k}-j}} a_{s,k,i}.$$

In the context of word pair learning, PPE was proposed to capture the spacing effect ([Cepeda et al., 2008](#)) – the phenomena that spaced out practice repetitions slow down learning but increase retention rates – by introducing a weighting scheme that considers the delay between individual practice sessions. PPE assumes a multiplicative relationship between the number of prior attempts  $a_{s,k}$  with a time variable  $T_k$ . The model features a learning rate parameter  $c$  and the forgetting rate is controlled by parameters  $x$ ,  $b$  and  $m$ . These four hyperparameters need to be set by the user. We define  $\Delta_{s,k,i}$  to be the real time passed since student  $s$ ’s  $i$ -th response to KC  $k$ . The PPE

prediction is defined as

$$p_{\text{PPE}}(a_{s,t+1} = 1 \mid q_{s,t+1}, \mathbf{x}_{s,1:t}) = \sigma \left( \sum_{k \in \text{KC}(q_{s,t+1})} \beta_k + \gamma_k (a_{s,k}^c T_k^{-d_t}) \right) \quad (5)$$

$$T_k = \left( \sum_{i=1}^{a_{s,k}} \Delta_{s,k,i}^{1-x} \right) \left( \sum_{j=1}^{a_{s,k}} \frac{1}{\Delta_{s,k,j}^{-x}} \right), \quad d_t = b + m \left( \frac{1}{a_{s,k}} \sum_{i=1}^{a_{s,k}} \frac{1}{\ln(\Delta_{s,k,i} - \Delta_{s,k,i+1} + e)} \right).$$

DAS3H is a more recent model which combines aspects of IRT and PFA and extends them with time-window based count features. It defines a set of time windows  $W = \{1/24, 1, 7, 30, +\infty\}$  measured in days. For each window  $w \in W$ , DAS3H determines the number of prior correct responses ( $c_{s,k,w}$ ) and overall attempts ( $a_{s,k,w}$ ) of student  $s$  on KC  $k$  which fall into the window. A scaling function  $\phi(x) = \log(1 + x)$  is applied to avoid features of large magnitude. The DAS3H prediction is defined as

$$p_{\text{DAS3H}}(a_{s,t+1} = 1 \mid q_{s,t+1}, \mathbf{x}_{s,1:t}) = \sigma \left( \alpha_s - \delta_{q_{s,t+1}} + \sum_{k \in \text{KC}(q_{s,t+1})} \beta_k + \sum_{k \in \text{KC}(q_{s,t+1})} \sum_{w=0}^{W-1} \theta_{k,2w+1} \phi(c_{s,k,w}) - \theta_{k,2w+2} \phi(a_{s,k,w}) \right). \quad (6)$$

[Gervet et al. 2020](#) performed a comparative evaluation of student performance modeling algorithms across 9 real-world datasets. They also evaluated the effects of question, KC and total count as well as time window based count features leading them to a new logistic regression model referred to as Best-LR. Best-LR is similar to DAS3H, but does not use time-window features and uses  $c_s$  and  $f_s$  as additional features that capture the total number of prior correct and incorrect responses. The Best-LR prediction is defined as

$$p_{\text{Best-LR}}(a_{s,t+1} = 1 \mid q_{s,t+1}, \mathbf{x}_{s,1:t}) = \sigma \left( \alpha_s - \delta_{q_{s,t+1}} + \phi(c_s) + \phi(f_s) + \sum_{k \in \text{KC}(q_{s,t+1})} \beta_k + \gamma_k \phi(c_{s,k}) + \rho_k \phi(f_{s,k}) \right). \quad (7)$$

Overall, we can see that these four logistic regression models are mainly based on two types of features: (i) one-hot encodings that allow the models to infer question and KC specific difficulty; (ii) count based features that summarize a student's past interaction history with the system computed at various level of granularity (total/KC/question-level) potentially augmented with time windows and time-based weighting schemes to introduce a temporal dimension in the predictions. Looking back at the dataset discussion provided in Section 3, we note that the current feature sets only use a small fraction of the information collected by the tutoring systems. In the following, we aim at increasing the amount of utilized information by exploring alternative types of features which can be extracted from the log data and can serve as foundation for future student performance modeling algorithms.

### 4.3 FEATURES BASED ON RICH OBSERVATIONAL DATA

Tutoring systems collect a wide variety of observational data during the learning process. Here we discuss a range of alternative features leveraging alternative types of log data. The individual



features can be used to augment the logistic regression models discussed in Subsection 4.2, but might also be combined with deep learning based modeling techniques. As shown by Table 2, each dataset captures different aspects of the learning process and supports a different subset of the discussed features.

**Temporal features:** Many performance modeling techniques treat student interactions as discrete tokens and omit timestamp information which can be indicators of cognitive processes such as short-term memorization and forgetting. DAS3H uses time window based count features to summarize the user history. Here we discuss two additional types of temporal features: (i) one-hot encoded datetime and (ii) the interaction time based features introduced by Shin et al. 2021. By providing the model with information about the specific week and month of interaction we try to capture effects of school work outside the ITS on student learning. The hour and day encodings aim at tracing temporal effects on a smaller timescale. For example, students might produce more incorrect responses when studying late at night. Recently, Shin et al. 2021 introduced a deep learning approach which employs *elapsed time* and *lag time* features to capture temporal aspects of student behavior during question solving activities. Elapsed time measures the time span from question display to response submission. The idea is that a faster response is correlated with student proficiency. Lag time measures the time passed between the completion of the previous exercise until the next question is received. Lag time can be indicative for short-term memorization and forgetting. The lag time between questions can be affected by student behavior and system design choices. Some examples are displayed explanations after incorrect responses, pop-up messages, and study breaks. For our experiments we convert elapsed time and lag time values,  $x$ , to scaled values  $\phi(x) = \log(1 + x)$  and also use the categorical one-hot encodings from Shin et al. 2021. There, elapsed time is capped off at 300 seconds and categorized based on the integer second. Lag time is rounded to integer minutes and assigned to one of 150 categories (0, 1, 2, 3, 4, 5, 10, 20, 30, . . . , 1440). We evaluate two variations of the features. In the first version we compute elapsed time and lag time values based on interactions with the *current* question. In the second version we compute them based on interactions with the *prior* question. Because it is unknown how long a student will take to answer a question before the question is asked, we cannot realistically use this elapsed time feature for predicting correctness in answering a new question. Therefore, we omit the elapsed time feature for the current question in our experiments described in Subsection 5.3.

**Learning context:** Knowledge about the context a learning activity is placed in can be informative for performance predictions. For example, all four datasets considered in this work group learning activities into distinct context categories we call study modules (e.g., pre-test, effective learning, review, . . .). Here effective learning study modules try to teach novel KCs, whereas review study modules aim at deepening proficiency related to KCs the student is already familiar with. Providing a student performance model with information about the corresponding study module can help it adapt its predictions to these different contexts. Additional context information is provided by the exercise structure. For example, the `ElemMath2021` dataset marks each interaction with a course and topic identifier and offers a manually curated question difficulty score. The `EdNet` dataset labels questions based on which part of the TOIEC<sup>®</sup> exam they address and groups them into bundles—a bundle is a set of questions which is asked together. Large parts of the `ElemMath2021` dataset were collected in physical learning centers where students can study under the supervision of human teachers and each interaction contains school and teacher identifiers. This information can be useful because students visiting the same learning center might exhibit similar strengths and weaknesses. In Section 5 we evaluate the

potential benefits of learning context information for student performance modeling by encoding the categorical information into one-hot vectors and passing it into logistic regression models. Additionally, we evaluate count features defined for the individual study modules and EdNet parts.

**Personal context:** By collecting a student’s personal attributes a tutoring system can offer a curriculum adapted to their individual needs. For example, information about a student’s age and grade can be used to select suitable learning materials. Further, knowledge about personal attributes can enable an ITS to serve as a vehicle for research on educational outcomes. For example, it is a well-studied phenomena that socioeconomic background correlates with educational attainment (see for example [White 1982](#), [Aikens and Barbarin 2008](#), [von Stumm 2017](#)). A related research question is how ITSs can be used to narrow the achievement gap ([Huang et al., 2016](#)). While both `ElemMath2021` and `EdNet` datasets indicate which modality students use to access the ITS, the `Eedi` dataset is the only one that provides more detailed personal information in form of age, gender, and socioeconomic status. Features extracted from these sensitive attributes need to be handled with care to avoid discrimination against individual groups of students. Outside of student performance modeling these features might be used to detect existing biases in the system. We evaluate one-hot encodings related to age, gender, and socioeconomic status.

**KC graph features:** In addition to a KC model three of the datasets provide a graph structure capturing semantic dependencies between the individual KCs and questions. We can leverage these graph structures by defining *prerequisite* and *post-requisite* count features. These features compute a student’s number of prior attempts and correct responses on KCs which are prerequisite and post-requisite to the KCs of the current question. The motivation for these features is that the mastery of related KCs can carry over to the current question. For example, it is very likely that a user that is proficient in multi-digit multiplication can solve single-digit multiplication exercises as well. The `Eedi` dataset provides a 4-level KC ontology tree which associates each question with a leaf. Here we compute a prerequisite feature by counting the number of interactions related to the direct parent node. We also investigate the use of sparse vectors taken from the KC graph adjacency matrix as an alternative way to incorporate the semantic information into our models.

**Learning material consumption:** Modern tutoring system offer a rich variety of learning materials in form of video lectures, audio recordings and written explanations. Information on how students interact with the available resources can benefit performance predictions. For example, watching a video lecture introducing a new KC can foster a student’s understanding before they enter the next exercise session. The `ElemMath2021` and `Ednet KT3` datasets record interactions with lecture video and written explanations. Count features capture the number of videos and texts a user has interacted with prior to the current question and can be computed on a per KC and an overall level. Video and reading time in minutes might also provide a useful signal. Another feature which can be indicative for proficiency is the number of videos a student skips. The `Junyi15` datasets provides us with information on hint usage. We experiment with features capturing the number of used hints as well as minutes spent on reading hints aggregated on a per KC and an overall level.

## 4.4 FEATURES BASED ON RESPONSE CORRECTNESS

### 4.4.1 Smoothed Average Correctness

The performance of logistic regression models relies on a feature set which offers a suitable representation of the problem domain. Unlike neural network based approaches, linear and linear logistic regression models are unable to recover higher-order dependencies between individual features on their own. For example, given two separate features describing a student’s number of prior correct responses,  $c$ , and overall attempts,  $a$ , the linear logistic model is unable to infer the ratio  $r = c/a$  of average correctness which can be a valuable indicator of a student’s long-term performance. To mitigate this particular issue, we introduce an additional feature,  $r_s$  capturing average correctness of student  $s$  over time as

$$r_s = \frac{c_s + \eta\bar{r}}{a_s + \eta} \quad (8)$$

where  $a_s$  is the number of questions attempted by student  $s$ , and  $c_s$  is the number of questions the student answered correctly. Here,  $\bar{r}$  is the average correctness rate over all other students in the dataset, and  $\eta \in \mathbb{N}$  is a smoothing parameter which biases the estimated average correctness rate,  $\bar{r}_s$  of student  $s$  towards this all-students average  $\bar{r}$ . The use of smoothing reduces the feature variance during a student’s initial interactions with the ITS. Related to our approach R-PFA (Galyardt and Goldin, 2015) introduced a feature that computes a recency-weighted proportion of correct responses for each individual KC via an exponential decay function. The LKT framework Pavlik et al. 2021 provides multiple features that can trace the ratio of correct student responses. By implementing multiple recency-based decay functions LKT extends R-PFA and introduces features that capture the overall correctness rate of each individual student. Unlike our formulation, LKT does not allow the use of a smoothing parameter to reduce the feature variance by biasing towards the all-student average. We calibrated the smoothing parameter for our experiments by evaluation  $\eta \in \{0, 1, 5, 10, 25, 50, 100, 250\}$  using the `ElemMath2021` dataset. Smoothing yielded benefits and we settled on  $\eta = 5$  for which we observed the largest accuracy improvements.

### 4.4.2 Response Patterns

Student performance modeling approaches that employ recurrent or transformer based neural networks take interaction sequences describing *multiple* student responses as input (e.g., the most recent 100). Provided this time-series data, deep learning algorithms are able to discover patterns and temporal dependencies in student behavior without requiring additional human feature engineering. Among the logistic regression models discussed in Subsection 4.2 only DAS3H incorporates temporal information in form of time window based count features to summarize student interactions over time scales from one hour to one months. While DAS3H can capture effects of cognitive processes such as long-term memorization and forgetting, it is still at disadvantage compared to deep learning based approaches that can also infer aspects of student behavior occurring on smaller time-scales. Indeed, it has been shown that a student’s most recent responses have a large effect on DKT performance predictions (Ding and Larson 2019, Ding and Larson 2021).

Here, inspired by the use of n-gram models in the NLP community (e.g., Manning and Schütze 1999), we propose *response patterns* as a feature which allows logistic regression models to infer aspects impacting short-term student performance. At time  $t$ , a response pattern  $r_t \in \mathbb{R}^{2^n}$

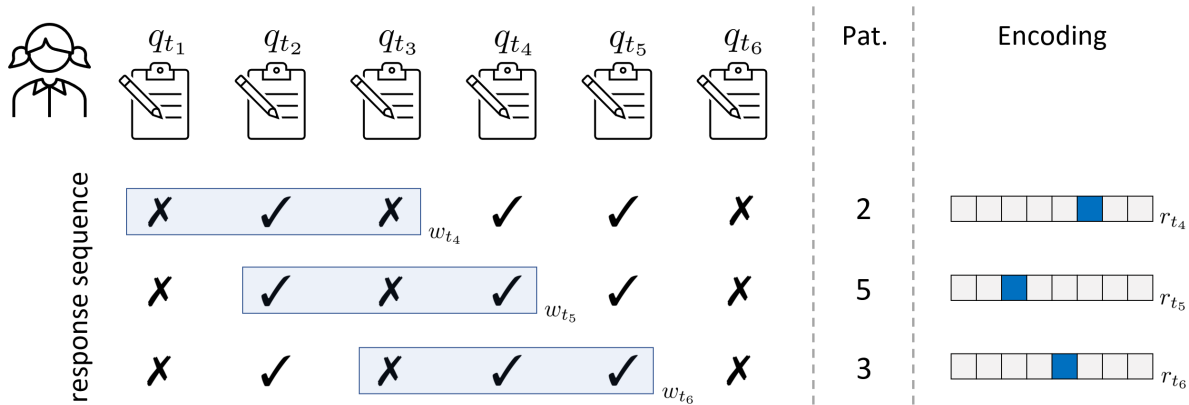


Figure 3: A response pattern  $r_t$  is a one-hot encoding vector which represents the binary sequence  $w_t$  formed by a student’s  $n$  most recent correct and incorrect responses. Logistic regression models can use response patterns to infer aspects of a student’s short-term behavior.

is defined as a one-hot encoded vector that represents a student’s sequence of  $n \in \mathbb{N}$  most recent responses  $w_t = (a_{t-n}, \dots, a_{t-1})$  formed by binary correctness indicators  $a_{t-n}, \dots, a_{t-1} \in \{0, 1\}$ . The encoding process is visualized by Figure 3. For our experiments we calibrated the pattern length by evaluating  $n \in \{1, 2, 3, \dots, 14\}$  using the `ElemMath2021` dataset. We settled on  $n = 10$  for which we observed the largest accuracy improvements. Response patterns are designed to allow logistic regression models to capture momentum in student performance. They allow the model to infer how challenging the current exercise session is for the user and can also be indicative for question skipping behavior. In Subsection 5.3 we combine response patterns, smoothed average correctness and DAS3H time window features to propose Best-LR+, a regression model which offers performance competitive to deep learning based techniques while only relying on information related to response correctness.

## 5 EXPERIMENTS

We evaluate the benefit of alternative features extracted from different types of log data for student performance modeling using the four large-scale datasets discussed in Section 3. After an initial feature evaluation, we combine helpful features to form novel state-of-the-art logistic regression models. Motivated by the size of the recent datasets, we also investigate two ways of partitioning the individual datasets to train multiple assessment models targeting different parts of the learning process. First, we show how partitioning a student’s interaction sequence by response number can be used to train *time-specialized* models – focusing on responses submitted earlier or later in the interaction sequence – to mitigate the cold start problem (Gervet et al. 2020, J. Zhang et al. 2021). We then analyze how features describing the learning context (e.g., study module, topic, course, ...) can be used to train multiple *context-specialized* models whose individual predictions can be combined to improve overall prediction quality even further.

### 5.1 EVALUATION METHODOLOGY

To be in line with prior work, we start data preparation by filtering out students with less than ten answered questions (Pandey and Karypis 2019, Choffin et al. 2019, Piech et al. 2015, Gervet




Figure 4: Accuracy, true positive rate and false positive rate of a Best-LR model trained on the ElemMath2021 dataset for different classification thresholds. The classification threshold converts the probability output by the model into a binary correct/incorrect prediction.




Figure 5: The receiver operating characteristic (ROC) curve captures the relationship between true and false positive rates as the classification threshold is swept from 0 to 1. The AUC score stands for the Area Under the Curve. A perfect system will have an AUC of 1.0.

et al. 2020). The Ednet<sub>KT3</sub> and Eedi dataset both contain questions annotated with multiple KCs which yields difficulties for some of the deep learning baselines (DKT and SAKT). In those cases we introduce new artificial KCs to represent each unique combination of original KCs. In all our experiments we perform a 5-fold cross-validation on the student level. In each fold 80% of the students are used for training and parameter selection and 20% are used for testing. Thus, all test results are obtained only from predictions over students who were not observed during model training. We report performance in terms of prediction accuracy (ACC) and area under curve (AUC). The AUC score captures the area under the receiver operating characteristic (ROC) curve and is a popular evaluation metric for student performance prediction algorithms. The ROC curve plots the true-positive rate against the false-positive rate at all decision thresholds. One way of interpreting the AUC score is viewing it as the probability of assigning a random correct student response a higher probability of correctness than a random incorrect response. As concrete examples we visualize the predictive performance and ROC curve for a Best-LR model trained on the ElemMath2021 dataset in Figures 4 and 5.

For the computation of average correctness and response pattern features we set smoothing parameter  $\eta = 5$  and sequence length  $n = 10$  respectively (discussed in Sections 4.4.1 and 4.4.2). For R-PFA we followed (Galyardt and Goldin, 2015) by fixing the number of ghost attempts  $g = 3$  and determined decay rate parameters  $d_F$  and  $d_R$  by selecting the best value in  $\{0.1, \dots, 0.9, 1.0\}$  for each dataset. For PPE we followed (Walsh et al., 2018) by fixing  $c = 0.1$  and  $x = 0.6$  and evaluated 20 equally-spaced decay parameters  $b \in [0.01, 0.05]$  and  $m \in [0.02, 0.04]$  for each dataset. Additional implementation details of the different features as well as ITS-specific considerations are provided in Appendix A. For the evaluation of the deep learning models, we performed a grid-search and defined hyperparameter search spaces which extend the hyperparameters selected in the cited references (DKT, (Piech et al., 2015); SAKT, (Pandey and Karypis, 2019); SAINT, (Choi et al., 2020); SAINT+, (Shin et al., 2021)). A detailed list of the used hyperparameter search spaces is provided in Appendix B. All models were trained for 100 epochs without learning rate decay. For our logistic regression experiments we rely on Scikit-learn (Pedregosa et al., 2011) and all deep learning architectures were implemented using



PyTorch (Paszke et al., 2019). Our implementation of the regression models uses combinations of attempt and correct count features which is a slight deviation from the original PFA and Best-LR formulations which count the number of prior correct and incorrect responses. While the individual features have a different interpretation, the two feature pairs are collinear to each other and provide identical information to the model. The code to reproduce our experiments as well as links to used datasets are shared on GitHub<sup>2</sup>.

## 5.2 UTILITY OF INDIVIDUAL FEATURES WITH LOGISTIC REGRESSION

To evaluate the effectiveness of the different features discussed in Section 4 we perform two experiments in the context of logistic regression modeling. In the first experiment we analyze each feature by training a logistic regression model using only that feature and a bias (constant) term. This allows us to quantify the predictive power of each feature in isolation. In the second experiment we evaluate the marginal utility gained from each feature by training an augmented Best-LR model (Eq. 7) based on the Best-LR feature set plus the additional feature. This is particularly interesting because the Best-LR feature set was created by combining features proposed by various prior regression models and it achieves state-of-the-art performance on multiple educational datasets (Gervet et al., 2020). If a feature can provide marginal utility on top of the Best-LR feature set it is likely that it can contribute to other student performance modeling techniques as well and should be captured by the ITS.

ACC and AUC scores achieved by the logistic regression models trained on the *individual features* in isolation are shown in Table 3. Somewhat unsurprisingly, the question (item) identifier – representing the main feature used by IRT (EQ 2) – is effective for all datasets and is the single most predictive feature for `ElemMath2021` and `EdNet_KT3`. KC identifiers are very informative as well, but to a lesser extent than the question identifiers. This indicates varying difficulty among questions targeting the same KC. The PFA model (EQ 3) which only estimates KC difficulties is unable to capture these variations on the question-level. The attempt and correct count features are informative for all datasets and there is benefit in using total-, KC- and question-level counts in combination. Question-level count features are informative for `Junyi15` which is the only dataset that frequently presents users previously encountered questions. Combined R-PFA’s two recency-weighted features deliver more accurate performance predictions than PPE’s spacing feature. DAS3H introduced count features computed for different time windows (TWs). TW based count features are an extension of the standard count features and yield additional benefits for all four settings. The elapsed time and lag time features are based on user interaction times and are both informative. The datetime features describing month, week, day and hour of interaction yield little signal on their own.

Looking at the individual learning context features we observe varying utility. The study module feature is informative for all datasets, as are the counts of student correct responses and attempts associated with different study modules. Predictions for the `Eedi` dataset benefit largely from information about the specific group a student belongs to. Reasons for this could be varying group proficiency levels and differences in diagnostic test difficulty. Bundle and part identifiers both provide a useful signal for `EdNet_KT3` and `Eedi`. Features describing personal information have little predictive power on their own. Among this family of features, social support yields the largest AUC score improvement over the always correct baseline. Features extracted from the prerequisite graph are effective for the three datasets that support them. The

---

<sup>2</sup><https://github.com/rschmucker/Large-Scale-Knowledge-Tracing>

Table 3: Individual feature performance. Each entry reports ACC and AUC scores achieved by a logistic regression model trained using only the corresponding feature and a bias term. The first row provides a baseline for comparison, where no features are considered and the model simply predicts the student will answer every question correctly. Dashed lines indicate this feature is not available in this dataset. Maximum ACC and AUC variances over the five-fold test data are 0.17% and 0.3% respectively.

Feature \ in %	ElemMath2021		EdNet KT3		Eedi		Junyi15	
	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC
always correct (baseline)	68.52	50.00	66.19	50.00	64.30	50.00	82.99	50.00
question ID	72.30	72.81	69.40	70.30	67.21	68.50	83.18	72.05
KC ID	69.64	66.63	66.29	58.87	64.29	57.46	83.02	64.06
total counts	70.14	64.86	66.43	59.54	69.03	72.00	83.21	64.62
KC counts	70.71	66.58	66.17	60.29	68.36	69.74	83.72	68.89
question counts	68.52	50.32	66.19	52.54	64.30	50.00	84.19	73.11
combined counts	72.05	71.05	66.65	62.63	70.28	73.94	84.35	74.57
total TW counts	70.64	65.96	66.48	60.62	70.42	74.32	83.60	68.27
KC TW counts	70.72	66.76	66.25	61.22	68.78	70.35	83.99	71.23
question TW counts	68.53	50.32	66.19	53.04	64.30	50.00	84.25	73.86
combined TW counts	72.30	71.66	66.79	64.09	71.11	75.18	84.39	75.41
R-PFA $F$	69.95	59.97	66.22	54.72	67.92	65.55	83.74	68.15
R-PFA $R$	68.57	63.56	66.18	60.28	64.30	61.94	83.37	72.71
R-PFA $F + R$	70.42	66.55	66.37	61.69	68.70	70.35	84.16	73.78
PPE Count	69.22	63.43	66.19	59.03	64.31	56.40	83.03	64.66
current elapsed time	69.72	61.50	66.18	57.36	-	-	83.22	66.23
current lag time	68.52	51.57	66.19	51.78	-	-	82.99	70.69
prior elapsed time	69.44	55.85	66.18	52.49	-	-	83.00	60.17
prior lag time	68.52	50.58	66.19	52.00	-	-	82.99	53.79
month	68.52	51.39	66.19	50.73	64.30	52.62	82.99	51.44
week	68.52	52.06	66.19	50.71	64.30	53.07	82.99	51.69
day	68.52	50.51	66.19	50.26	64.30	51.55	82.99	51.21
hour	68.52	50.36	66.19	50.62	64.30	51.63	82.99	52.13
study module ID	68.52	55.28	66.29	54.44	64.64	58.34	82.99	66.40
study module counts	70.35	62.76	66.18	54.00	68.13	68.82	83.17	65.85
teacher/group	68.37	56.52	-	-	66.85	66.84	-	-
school	68.50	55.62	-	-	-	-	-	-
course	68.53	54.56	-	-	-	-	-	-
topic	68.54	59.27	-	-	-	-	-	-
difficulty	68.55	54.73	-	-	-	-	-	-
bundle/quiz	-	-	68.70	68.21	65.05	62.99	-	-
part/area ID	-	-	66.19	56.21	-	-	83.02	57.06
part/area counts	-	-	66.53	61.01	-	-	83.48	66.16
age	-	-	-	-	64.30	53.30	-	-
gender	-	-	-	-	64.30	51.51	-	-
social support	-	-	-	-	64.30	55.13	-	-
platform	68.52	50.17	66.19	51.73	-	-	-	-
prereq IDs	69.64	66.63	-	-	64.29	57.46	83.18	72.05
prereq counts	71.73	69.72	-	-	69.25	71.88	84.42	75.97
postreq IDs	69.63	66.63	-	-	-	-	83.18	72.05
postreq counts	71.12	69.25	-	-	-	-	84.34	76.03
videos watched counts	68.52	53.44	66.19	54.24	-	-	-	-
videos skipped counts	68.59	57.23	66.19	53.27	-	-	-	-
videos watched time	68.52	53.29	66.19	54.05	-	-	-	-
reading counts	68.90	58.92	66.19	55.28	-	-	-	-
reading time	68.50	54.92	66.19	51.48	-	-	-	-
hint counts	-	-	-	-	-	-	82.99	60.00
hint time	-	-	-	-	-	-	82.98	59.73
smoothed avg correct	70.18	65.14	66.49	59.89	69.16	72.10	82.99	66.40
response pattern	70.86	64.98	66.39	59.75	69.67	72.71	84.02	70.51

logistic regression models trained on prerequisite and post-requisite count features exhibit the best AUC scores on Junyi15. The features describing study material consumption all provide a predictive signal, but do not yield enough information for accurate performance predictions on their own. The smoothed average correctness and response pattern features lead to good performance on all four datasets.

Moving forward we evaluate the marginal utility of each feature by evaluating the performance of logistic regression models trained using the *Best-LR feature set* augmented with this particular feature. Table 4 shows the results of this experiment. The combination of the two R-PFA features is beneficial for all four datasets and increases the AUC score for Junyi15 by 1.75%. The PPE feature is also helpful, but to a lesser extent than the R-PFA features. Time window based count features lead to improved predictions for all datasets. In combination, they improve the AUC scores of EdNet\_KT3 and Junyi by 1.29% and 2.2% respectively. The elapsed time and lag time features offer a useful signal for the three datasets which capture temporal interaction data. The datetime features provide little utility in all cases (the best observed value is a 0.08% AUC improvement for Junyi15).

Context information about the study module a question is placed in is beneficial in all settings and leads to the largest AUC score improvement over the Best-LR baseline model for Junyi15. Indication of the group a student belongs to improves model performance for Eedi. Knowing which topic a student is currently visiting is a useful signal for ElemMath2021. Information related to bundle and part structure improves EdNet\_KT3 and Eedi predictions. ElemMath2021's manually assigned difficulty ratings lead to no substantial improvements likely due to the fact that the Best-LR feature set allows the model to learn question difficulty on its own. The features describing a student's personal information provide little marginal utility. Count features derived from the question/KC prerequisite graphs yield a sizeable improvement in assessment quality. Features targeting study material consumption yield some marginal utility when available. Discrete learning material consumption counts lead to larger improvements than the continuous time features. The smoothed average correctness feature leads to noticeable improvements for all four datasets. The response pattern feature enhances assessments in all cases and yields the largest improvements for the ElemMath2021 and Eedi dataset (over 0.5% AUC). This shows that the last few most recent student responses yield a valuable signal for performance prediction.

Overall, we observe that a variety of alternative features derived from different types of log data can enhance student performance modeling. Tutoring systems that track temporal aspects of student behavior in detail can employ elapsed time and lag time features. Additional information about the current learning context is valuable and should be captured. While the study module features improve predictions for all datasets, other ITS-specific context features such as group, topic and bundle identifiers vary in utility. The count features derived from the provided prerequisite and hierarchical graph structures increased prediction quality in all cases. Log data related to study material consumption also provides a useful signal. Beyond student performance modeling, this type of log data might also prove itself useful for evaluating the effects of individual video lectures and written explanations in future work. Lastly, the smoothed average correctness and response pattern features that only require answer correctness improve predictions for all four tutoring systems substantially.

Table 4: Augmented Best-LR performance. Each entry reports average ACC and AUC scores achieved by a logistic regression model trained using the Best-LR feature set augmented with a single feature. Maximum ACC and AUC variances over the five-fold test data are 0.15% and 0.13% respectively. The marker  $\times$  is used to indicate features that are used for the AugmentedLR assessment models described later in the paper. Dashed lines indicate this feature is not available in this dataset.

Feature \ in %	ElemMath2021		EdNet KT3		Eedi		Junyi15					
	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC				
Best-LR (baseline)	75.69	78.44	$\times$	70.69	72.94	$\times$	73.43	79.01	$\times$	84.25	76.20	$\times$
question counts	75.70	78.45		71.24	73.53		73.43	79.01		84.74	77.76	
total TW counts	75.90	78.78		70.75	73.07		73.94	79.73		84.37	76.68	
KC TW counts	75.72	78.54		70.94	73.27		73.64	79.30		84.47	76.98	
question TW counts	75.71	78.47		71.41	74.05		73.43	79.01		84.78	78.27	
combined TW counts	75.95	78.88	$\times$	71.51	74.23	$\times$	73.99	79.79	$\times$	84.84	78.40	$\times$
R-PFA $F$	75.73	78.50		70.77	73.14		73.71	79.36		84.56	77.23	
R-PFA $R$	75.72	78.49		70.75	73.10		73.57	79.15		84.72	77.88	
R-PFA $F + R$	75.74	78.51	$\times$	70.80	73.25	$\times$	73.73	79.38	$\times$	84.77	77.95	$\times$
PPE Count	75.73	78.53	$\times$	70.73	73.07	$\times$	73.44	79.02	$\times$	84.34	76.75	$\times$
current elapsed time	76.07	78.97		71.01	74.14		-	-		84.38	77.62	
current lag time	75.79	78.54	$\times$	70.71	73.01	$\times$	-	-		84.34	76.65	$\times$
prior elapsed time	75.88	78.67	$\times$	70.77	73.11	$\times$	-	-		84.26	76.42	$\times$
prior lag time	75.76	78.52		70.75	73.10		-	-		84.26	76.45	
month	75.70	78.45		70.69	72.95		73.43	79.02		84.25	76.21	
week	75.70	78.45		70.69	72.95		73.44	79.02		84.26	76.21	
day	75.70	78.45		70.69	72.94		73.43	79.02		84.26	76.24	
hour	75.70	78.45		70.69	72.94		73.44	79.01		84.27	76.28	$\times$
study module ID	75.77	78.60	$\times$	71.40	73.88	$\times$	73.52	79.10	$\times$	84.56	82.39	$\times$
study module counts	75.74	78.54		70.76	73.01		73.52	79.10		84.27	76.30	
teacher/group	75.68	78.39		-	-		74.00	79.63	$\times$	-	-	
school	75.72	78.48		-	-		-	-		-	-	
course	75.72	78.49		-	-		-	-		-	-	
topic	75.74	78.53	$\times$	-	-		-	-		-	-	
difficulty	75.70	78.45		-	-		-	-		-	-	
bundle/quiz ID	-	-		70.69	72.94		73.76	79.43	$\times$	-	-	
part/area ID	-	-		70.69	72.94		-	-		84.25	76.20	
part/area counts	-	-		70.73	73.05	$\times$	-	-		84.26	76.23	
age	-	-		-	-		73.45	79.03		-	-	
gender	-	-		-	-		73.43	79.01		-	-	
social support	-	-		-	-		73.44	79.02		-	-	
platform	75.70	78.45		70.68	72.94		-	-		-	-	
prereq IDs	75.70	78.45		-	-		73.43	79.01		84.25	76.20	
prereq counts	75.91	78.77	$\times$	-	-		73.54	79.15	$\times$	84.91	78.20	$\times$
postreq IDs	75.69	78.45		-	-		-	-		84.25	76.20	
postreq counts	75.81	78.64	$\times$	-	-		-	-		84.83	78.02	$\times$
videos watched counts	75.75	78.51	$\times$	70.70	73.04	$\times$	-	-		-	-	
videos skipped counts	75.72	78.49		70.70	73.00	$\times$	-	-		-	-	
videos watched time	75.72	78.48		70.68	72.99		-	-		-	-	
reading counts	75.75	78.58	$\times$	70.69	72.95		-	-		-	-	
reading time	75.70	78.45		70.69	72.96		-	-		-	-	
hint counts	-	-		-	-		-	-		84.31	76.59	$\times$
hint time	-	-		-	-		-	-		84.27	76.40	$\times$
smoothed avg correct	75.78	78.62	$\times$	70.81	73.22	$\times$	73.49	79.13	$\times$	84.28	76.42	$\times$
response pattern	76.03	78.99	$\times$	70.82	73.24	$\times$	74.32	80.10	$\times$	84.72	77.65	$\times$

### 5.3 INTEGRATING FEATURES INTO LOGISTIC REGRESSION

We have seen how individual features can be integrated into the Best-LR model to increase prediction quality. We now experiment with augmented logistic regression models that incorporate combinations of *multiple* beneficial features. The number of parameters learned by the different

approaches is provided in Table 6. Because each dataset provides different types of log data it only supports a subset of the explored features. The first model we propose and evaluate is Best-LR+.

**Best-LR+ method.** This method augments the Best-LR method (the current state-of-the-art logistic regression method (EQ 7)) by adding R-PFA's two recency-weighted features, PPE's spacing time feature and DAS3H's time window-based count features on total-, KC- and question-level as well as smoothed average correctness and response pattern features. All of these features can be calculated using only the raw question-response information from the student log data.

Note that because the features used in Best-LR+ rely only on question-response log data, this method can be applied to all four datasets, as well as many earlier datasets that lack new features available in the four we study here. We define Best-LR+ in this way, to explore whether tutoring systems that log only this question-response data can improve the accuracy of the student performance modeling by calculating and adding in these count features.

We further propose dataset specific augmented logistic regression models (AugmentedLR) by combining the helpful features marked in Table 4.

**AugmentedLR method.** This logistic regression method uses all of the features employed in the Best-LR model and adds in all of the marked features from Table 4, which were found to individually provide AUC score improvements of more than 0.05% over the state-of-the-art results of Best-LR.

Note that the AugmentedLR method employs a superset of the features used by Best-LR+. Because the features used in AugmentedLR go beyond simple logs of questions and responses, this method can only be applied to datasets that capture this additional information. We define AugmentedLR in this way, to explore whether ITS systems that do not yet log these augmented features can improve the accuracy of their student performance modeling by capturing and utilizing these additional features in their student log data. Note that some of the features marked in Table 4 and used by AugmentedLR are available in only a subset of our four datasets. During our AugmentedLR feature selection we made two additions to the over 0.05% AUC improvement rule to mitigate redundancy: (i) In cases where count and ID feature target the same attribute, we select the one yielding the larger AUC improvements. (ii) Information about the current lag time is preferred over prior lag time.

Experimental results for these two proposed logistic regression models are presented in Table 5, along with results for previously published logistic regression approaches IRT, PFA, R-PFA, PPE, DAS3H, Best-LR and previously published deep learning approaches DKT, SAKT, SAINT and SAINT+.

Examining the results, first compare the results for Best-LR+ to the previous state of the art in logistic regression methods, Best-LR. Notice that Best-LR+ benefits from its additional question-response features and outperforms Best-LR in all four datasets. On EdNet KT3 and Juni15 Best-LR+ improves the AUC scores of the previous best logistic regression models (Best-LR) by more than 1.7%. These results suggest that the additional features used by Best-LR+ should be incorporated into student performance modeling in any ITS, even those for which the student log data contains only the sequence of question-response pairs.

Examining the results for AugmentedLR, it is apparent that this provides the most accurate method for student modeling across these four diverse datasets, among the logistic regression



Table 5: Comparative evaluation of student performance modeling algorithms across four large-scale datasets. The first six table rows correspond to previously studied logistic regression methods, the next four to previously studied deep neural network approaches, and the final two rows correspond to the two new logistic regression methods introduced in this paper. Maximum ACC and AUC variances over the five-fold test data are 0.16% and 0.17% respectively. Because the `Eedi` dataset does not provide response time information it does not accommodate SAINT+ and there is no corresponding entry.

Model \ in %	ElemMath2021		EdNet	KT3	Eedi		Junyi15	
	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC
IRT	72.30	72.81	69.40	70.30	67.21	68.50	83.18	72.05
PFA	71.68	70.80	66.48	61.95	68.44	70.28	83.85	70.19
R-PFA	71.60	70.80	66.60	62.92	68.90	71.01	84.35	74.13
PPE	70.08	67.59	66.51	59.93	64.32	57.95	83.09	64.81
DAS3H	74.08	75.82	70.31	72.16	71.64	76.13	84.43	76.73
Best-LR	75.69	78.44	70.69	72.94	73.43	79.01	84.25	76.20
DKT	76.46	79.71	71.77	74.92	<b>75.41</b>	<b>81.55</b>	85.50	80.62
SAKT	75.90	78.44	71.53	74.11	74.51	80.31	85.16	79.59
SAINT	75.87	77.88	71.40	73.69	74.56	80.38	85.10	79.51
SAINT+	76.04	78.15	71.54	73.94	-	-	85.18	79.71
Best-LR+	76.23	79.35	71.69	74.65	74.55	80.40	85.05	79.12
AugmentedLR	<b>76.59</b>	<b>79.87</b>	<b>71.89</b>	<b>75.00</b>	74.96	80.96	<b>86.35</b>	<b>86.03</b>

methods considered here. Even when considering recent deep learning based models, the AugmentedLR models outperform all logistic regression and deep neural network methods on all datasets except `Eedi` where the deep network method DKT leads to the best performance predictions. Especially on the `Junyi15` dataset AugmentedLR increases ACC by 2.1% and AUC by over 9.8%, compared to Best-LR. These results strongly suggest that all ITSs might benefit from logging the alternative features of AugmentedLR to improve the accuracy of their student performance modeling.

#### 5.4 THE COLD START PROBLEM

Student performance modeling techniques use a student’s interaction history to make predictions about their ability to solve different problems over time. When a new student starts using the ITS, little or no interaction history is yet available, resulting in a new student “cold start” problem for estimating the performance of new students. Most student performance models therefore require a burn-in period of student use of the ITS before they can accurately estimate student performance (Gervet et al. 2020, J. Zhang et al. 2021). Here, to ensure a gratifying user experience and to improve early on retention rates, we show how one can mitigate the cold start problem by training multiple *time-specialized* assessment models. We start by splitting the question-response sequence of each student into multiple distinct partitions based on their ordinal position in the student’s learning process (i.e., partition *50-100* will contain the 50th to 100th response of each student). We then train a separate time-specialized model for each individual partition. The motivation for this is that the way observational data needs to be interpreted can change over time. For example, during the beginning of the learning process one might put more focus on question and KC identifiers, while later on count features provide a richer signal. With this approach a

Table 6: Number of parameters learned by different student performance models. The number of parameters is heavily dependent on the number of questions and KCs used by the underlying ITS.

	ElemMath2021	EdNet KT3	Eedi	Junyi15
# of unique questions	59,892	13,169	27,613	835
# of KCs	4,191	293	388	41
IRT	59,571	11,556	27,614	723
PFA	12,574	904	1,165	124
R-PFA	12,574	904	1,165	124
PPE	8,383	603	777	83
DAS3H	105,672	14,867	31,882	1,174
Best-LR	72,146	12,461	28,780	848
DKT	21,529,751	526,526	10,676,751	184,801
SAKT	6,730,901	1,033,051	3,191,301	125,101
SAINT	16,727,873	8,026,049	7,666,497	1,618,049
SAINT+	4,194,241	1,349,889	-	148,865
Best-LR+	119,291	16,816	34,092	2,343
AugmentedLR	137,602	17,297	64,076	6,170
Combined-AugmentedLR	10,595,355	242,159	4,164,941	86,381

student’s proficiency is evaluated by different specialized models depending on the length of their prior interaction history.

We evaluate the technique of training multiple time-specialized logistic regression models using the Best-LR and Best-LR+ feature sets using the four educational datasets (additional results for AugmentedLR are provided by Table 8). In particular we induce the partitions using the splitting points  $\{0, 10, 50, 100, 250, 500, \infty\}$ . Figure 6 visualizes the results of a 5-fold cross-validation and compares predictive performance with a single *generalist* Best-LR model trained using all available data. Figure 7 shows the corresponding experiment for Best-LR+. For both Best-LR and Best-LR+ feature sets the use of time-specialized assessment models substantially improves prediction accuracy early on (i.e., during their first 10 question-response pairs) and mitigates the cold-start problem successfully for all four datasets. For EdNet KT3 and Junyi we observe AUC improvements of over 2% and 0.8% in early on predictions compared to the generalist Best-LR and Best-LR+ models respectively. Also, for the same two datasets the overall Best-LR performance (shown in Table 7) achieved by combining the predictions of the individual time-specialized models yields an over 0.5% increase in overall AUC scores and the time-specialized models consistently outperform the generalist model in each individual partition. For ElemMath2021 and Eedi datasets the time-specialized Best-LR models do not outperform the baseline consistently in each partition, but we still observe minor gains in predictive performance overall. Looking at the overall performance achieved by time-specialized Best-LR+ (Figure 7) and time-specialized AugmentedLR models (Table 8) we observe mixed results. While we observe consistent benefits for EdNet KT3 and Junyi15 the time-specialized models sometimes harm overall performance for ElemMath2021 and Eedi. This might be due to the increased data intensity of the Best-LR+ and AugmentedLR models which we will discuss in more detail in Section 5.5.




Figure 6: Comparison in AUC performance between a single Best-LR model trained using the entire training set (blue), and a collection of time-specialized Best-LR models trained on partitions of the response sequences (red). The horizontal axis shows the subset of question-response pairs used to train the specialized model. For example, the leftmost red point in each plot shows the AUC performance of a specialized model trained using only the first 10 question-response pairs of each student, then tested on the first 10 responses of other students. In contrast, the leftmost blue point shows the performance of the Best-LR model trained using all available data, and still tested on the first 10 responses of other students. The time-specialized models are able to mitigate the cold start problem for all four tutoring systems. On the EdNet KT3 and Junyi15 datasets the combined predictions of the specialized models increase overall performance substantially.

## 5.5 COMBINING MULTIPLE SPECIALIZED MODELS

The large-scale datasets discussed in this paper are aggregations of learning trajectories collected from users of varying ages and grades who complete a range of different courses. The internal heterogeneity of the datasets paired with the large number of recorded interactions naturally leads to the question of whether it is more advantageous to pursue a monolithic or a composite modeling approach. Conventional student performance modeling techniques follow a monolithic approach in which they train a single model using all available training data, but the results in the previous section show that with sufficiently large datasets it can be useful to partition the training data to train multiple time-specialized models. The idea of learning different classification rules for different parts of the dataset is a popular technique in the machine learning literature and a core principle behind decision tree algorithms (Breiman, 2001). A related EDM question is whether it is more beneficial to model KCs in separation or in combination. For example, while BKT partitions the student interaction sequences by KC to infer KC-specific knowledge estimates, DKT takes in entire sequences to output a single vector which describes student proficiency for all KCs. Starting from this observation Montero et al. 2018 explore a modified DKT approach called DKT-SM-SS. Analogous to BKT, DKT-SM-SS partitions the student interaction sequences by




Figure 7: Comparison in AUC performance between a single Best-LR+ model trained using the entire training set (blue), and a collection of time-specialized Best-LR+ models trained on partitions of the response sequences (red). The horizontal axis shows the subset of question-response pairs used to train the specialized model. For example, the leftmost red point in each plot shows the AUC performance of a specialized model trained using only the first 10 question-response pairs of each student, then tested on the first 10 responses of other students. In contrast, the leftmost blue point shows the performance of the Best-LR+ model trained using all available data, and still tested on the first 10 responses of other students. The time-specialized models are able to mitigate the cold start problem for all four tutoring systems. On the EdNet KT3 and Junyi15 datasets the combined specialized predictions also improve overall performance.

KC and trains a separate KC-specific DKT model for each partition. Comparing the performance of DKT-SM-SS with conventional DKT they found that DKT benefits from modeling the KCs in combination. When partitioning the dataset per KC our logistic regression modeling approach has an advantage over DKT-SM-SS in that the underlying feature set still captures information about interactions with all KCs.

In this section we explore the potential benefits of (1) training specialized models for specific questions, KCs, study modules, etc., and (2) combining the predictions from several of these models to obtain a final prediction of student performance. The motivation for considering partitioning the data to train models specialized to these different contexts is that it has the potential to recover finer nuances in the training data, just as the time-specialized models in the previous section did. For example, two algebra courses might teach the same topics, but employ different analogies to promote conceptual understanding, allowing students to solve certain types of questions more easily than others. Training a separate specialized model for each course can allow the trained models to capture these differences and improve overall assessment quality.

Table 7 shows performance metrics for sets of logistic regression models trained using the Best-LR features. As baselines we use a single Best-LR model trained on the entire dataset as well as a set of time-specialized Best-LR models (described in Subsection 5.4). Note that even

Table 7: Composite Best-LR performance for different partitioning schemes. The first two rows of the table give as baseline scores the ACC and AUC for the previously discussed Best-LR model and time-specialized Best-LR models (Subsection 5.4). The next 10 rows show results when training specialized models that partition the data based on single features such as question ID, KC ID, etc.. The final row shows the result of combining several of these models by taking a weighted vote of their predictions as described in the text. The marker **X** indicates the inputs used for this combination model for each dataset. Maximum ACC and AUC variances over the five-fold test data are 0.14% and 0.12% respectively.

Feature \ in %	ElemMath2021		EdNet KT3		Eedi		Junyi15				
	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC			
Best-LR	75.69	78.45	70.69	72.94	<b>X</b>	73.43	79.01	84.25	76.20	<b>X</b>	
Best-LR (time-spec.)	75.71	78.47	71.04	73.53	<b>X</b>	73.51	79.08	<b>X</b>	84.33	76.81	<b>X</b>
question ID specific	75.73	78.31	70.90	73.18		73.69	79.12	<b>X</b>	84.37	76.76	<b>X</b>
KC ID specific	75.81	78.63	<b>X</b>	70.80	73.11	<b>X</b>	73.50	79.08	84.26	76.26	
study module specific	75.96	78.94	<b>X</b>	71.65	74.43	<b>X</b>	73.73	79.30	84.69	82.67	<b>X</b>
teacher/group specific	70.62	66.91	-	-	-	72.99	78.15	-	-	-	
school specific	72.17	71.33	-	-	-	-	-	-	-	-	
course specific	75.79	78.59	<b>X</b>	-	-	-	-	-	-	-	
topic specific	75.82	78.68	-	-	-	-	-	-	-	-	
bundle/quiz specific	-	-	70.92	73.24	<b>X</b>	73.95	79.49	<b>X</b>	-	-	
part/area specific	-	-	70.71	73.01		-	-		84.25	76.23	<b>X</b>
platform specific	75.67	78.39	<b>X</b>	70.71	72.98	<b>X</b>	-	-	-	-	
Combined-Best-LR	76.08	79.13	71.75	74.65		74.13	79.83		84.71	82.90	

Table 8: Composite AugmentedLR performance for different partitioning schemes. The first two rows of the table give as baseline scores the ACC and AUC for the previously discussed AugmentedLR model and time-specialized AugmentedLR models (Subsection 5.4). The next 10 rows show results when training specialized models that partition the data based on single features such as question ID, KC ID, etc.. The final row shows the result of combining several of these models by taking a weighted vote of their predictions as described in the text. The marker **X** indicates the inputs used for this combination model for each dataset. Maximum ACC and AUC variances over the five-fold test data are 0.19% and 0.14% respectively.

Feature \ in %	ElemMath2021		EdNet KT3		Eedi		Junyi15					
	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC				
AugmentedLR	76.59	79.87	71.89	74.99	74.96	80.96	86.37	86.03				
AugmentedLR (time-spec.)	76.41	79.51	71.92	75.01	<b>X</b>	74.83	80.80	<b>X</b>	86.38	86.16	<b>X</b>	
question ID specific	75.09	77.37	70.16	72.36	74.33	80.14	86.12	85.72				
KC ID specific	75.73	78.49	71.21	73.86	74.59	80.57	86.16	85.82				
study module specific	76.71	80.03	<b>X</b>	72.02	75.34	<b>X</b>	74.92	80.93	<b>X</b>	86.42	86.19	<b>X</b>
teacher/group specific	66.96	61.16	-	-	72.53	77.62	-	-				
school specific	70.03	66.82	-	-	-	-	-	-				
course specific	76.49	79.69	<b>X</b>	-	-	-	-	-				
topic specific	75.98	78.89	-	-	-	-	-	-				
bundle/quiz specific	-	-	70.22	72.44	73.96	79.65	-	-				
part/area specific	-	-	71.91	75.09	-	-	86.29	85.99				
platform	76.55	79.78	71.92	75.04	-	-	-	-				
Combined-AugmentedLR	76.76	80.16	72.11	75.48	75.04	81.11	86.46	86.34				

though question and KC identifiers are already part of the Best-LR feature set they can still be effective splitting criteria. As shown in the table, training question specialized models improves predictive performance for all datasets except ElemMath2021 and KC specialized models are beneficial for all four datasets. Partitioning the data on the value of the study module feature and training specialized models for each study module yields the greatest improvements for

Table 9: Average and median number of training examples (i.e., question responses) available for different partitioning schemes. For example, when a specialized model is trained for each different question ID in the `ElemMath2021` dataset, the average number of training examples available per question-specific model is 393.

Feature	ElemMath2021		EdNet KT3		Eedi		Junyi15	
	Avg	Med	Avg	Med	Avg	Med	Avg	Med
question ID specific	393	118	1,445	930	718	351	35,207	8,380
KC ID specific	5,801	2,031	11,694	2,956	18,485	613	635,487	200,284
study module specific	3,901,401	2,487,865	2,385,616	810,120	336,183	29,589	5,083,895	3,421,938
teacher/group specific	1,370	297	-	-	1,674	554	-	-
school specific	11,992	5,620	-	-	-	-	-	-
course specific	329,695	69,125	-	-	-	-	-	-
topic specific	21,614	4809	-	-	-	-	-	-
bundle/quiz specific	-	-	1,970	1,265	1,146	120	-	-
part/area specific	-	-	2,385,616	1,344,293	-	-	2,824,386	605,666
platform	11,704,204	11,704,204	8,349,656	8,349,656	-	-	-	-

`ElemMath2021`, `EdNet KT3` and `Junyi15`, and also leads to substantial improvements on `Eedi`. It is interesting that training multiple specialized models for different study modules is more effective than augmenting the Best-LR feature set directly with the study module feature (Table 4). Topic and course specific data partitions improve performance predictions for the `ElemMath2021` dataset. On the other hand, school and teacher specific splits are detrimental and we observe large overfitting to the training data. Splits on the bundle/quiz level are effective for `EdNet KT3` and `Eedi`. While `Eedi` benefits from incorporating the group identifiers into the Best-LR model, group specialized models harm overall performance.

While fitting the question specific models on `ElemMath2021`, we observed severe overfitting behavior, where accuracy on the training data is much higher than on the test set. This is likely caused by the fact that `ElemMath2021` contains the smallest number of responses per question among the four datasets. Table 9 provides information about the average and median number of responses per model for the different data partitioning schemes.

We repeat the same experiment, but this time train logistic regression models using the ITS-specific AugmentedLR feature sets. Performance metrics are provided by Table 8. Unlike when using the Best-LR features, the question and KC specific AugmentedLR models have lower overall prediction quality than the original AugmentedLR model. These specialized AugmentedLR models contain many more parameters than their Best-LR counterparts due to the fact that they use more features to describe the student history (Table 6). This larger number of parameters requires a larger training dataset, which makes them more prone to overfitting. Still, splits on the study module, course, and part features yield benefits on multiple datasets. Even though the AugmentedLR models benefit less from training specialized models, these specialized models still exhibit higher performance than the ones based on the Best-LR feature set.

Finally, we discuss combining the strengths of the different models described in this section by combining their predictions into a final group prediction via a machine learning technique called Stacking (Wolpert, 1992). We do so by training a higher-order logistic regression model that takes as input the predictions of the different models, and outputs the final predicted probability that the student will answer the question correctly. The learned weight parameters in this higher-order logistic regression model essentially create a weighted voting scheme that combines the predictions of the different models. We determined which models to include in this group prediction by evaluating all different combinations using the first data split of the 5-fold validation



scheme and selecting the one yielding the highest AUC. The results of the best-performing combination models are shown in the final rows of Tables Table 7 (Combined-Best-LR) and Table 8 (Combined-AugmentedLR). Both tables also mark which models were used for this combination model, for each of the four datasets.

The combination models outperform the baselines as well as each individual set of specialized models for both Best-LR and AugmentedLR feature sets on all datasets. These combination models are the most accurate of all the logistic regression models discussed in this paper, and they also outperform all the deep learning baselines (Table 5) on all datasets except on Eedi where DKT is the only model that produces more accurate predictions. Looking at the Best-LR based combination models (Combined-Best-LR) we observe large AUC improvements of more than 0.65% over the baseline models. For EdNet KT3 and Junyi15 there is an increase in AUC scores of 1.12% and 6.09% respectively. The minimum number of individual predictions used by a combination model is 3 (for Eedi) and the maximum number is 6 (for EdNet KT3). For the AugmentedLR based combination models (Combined-AugmentedLR) we observe AUC score improvements between 0.15% (for Eedi) and 0.47% (for EdNet KT3) compared to the baseline models. The AugmentedLR models already contain many of the features used for partitioning and are more data intensive than the Best-LR based models which is likely to be the reason for the smaller performance increment. The predictions of the combined AugmentedLR models rely mainly on the predictions of time- and study module-specialized models. Only the combination model for ElemMath2021 uses course- instead of time-specialized models as its second signal.

We note that while we were able to enhance prediction quality using only simple single feature partitioning schemes, future work on better strategies to partition the training data to train specialized models is likely to yield even larger benefits.

## 6 RESULTS SUMMARY AND DISCUSSION

The main results of this paper show that the state of the art of student performance modeling can be advanced through new machine learning approaches, yielding more accurate student assessments and in particular more accurate predictions of which questions a student will be able to answer correctly at any given point in time as they move through the online course. In particular we show:

- State-of-the-art logistic regression approaches to student performance modeling can be further improved by incorporating a set of *new features that can be easily calculated from the question-response pairs* that appear in student log data from nearly all ITSs. For example, these include counts and smoothed ratios of correctly answered questions and overall attempts, partitioned by question, by knowledge component, and by time window, as well as specific sequences of correct/incorrect responses over the most recent student responses. We refer to the previous state-of-the-art logistic regression model as Best-LR (Gervet et al., 2020) and to the logistic regression model that incorporates these additional features as Best-LR+. Our experiments show that Best-LR+ yields more accurate student modeling than Best-LR across all four diverse ITS logs considered in this paper. We conclude that most tutoring systems that perform student modeling should benefit by incorporating these features.

- A second way of improving over the state of the art in student modeling is to incorporate *new types of features that go beyond the traditional question-response data* typically logged in all ITSs. For example, accuracy is improved by incorporating features such as the time students took to answer the previous questions, student performance on earlier questions associated with prerequisites to the knowledge component of the current question, and information about the study module (e.g., does the question appear in a pre-test, post-test, or as a practice problem). We conclude that future tutoring systems should log the information needed to provide these features to their student performance modeling algorithms.
- A third way to improve on the state of the art is to train *multiple, specialized student performance models* and then combine their predictions to form a final group prediction. For example, we found that training distinct logistic regression models on different partitions of the data (e.g., partitioning the data by its position in the sequential log, or by the knowledge component being tested) leads to improved accuracy. Furthermore, combining the predictions of different specialized models leads to additional accuracy improvements (e.g., combining the predictions of specialized models trained on different question bundles, with predictions of specialized models trained on different periods of time in the sequential log). We conclude that time-specialized models can help ameliorate the problem of assessing new students who have not yet created a long sequence of log data. Furthermore, we feel that as future tutoring systems are adopted by more and more students, the increasing size of student log datasets will make this approach of training and combining specialized models increasingly effective.
- Although our primary focus here is on logistic regression models, we also considered top-performing neural network approaches including DKT (Piech et al., 2015), SAKT (Pandey and Karypis, 2019), SAINT (Choi et al., 2020) and SAINT+ (Shin et al., 2021) as additional state-of-the-art systems against which we compare. Our experiments show that among these neural network approaches, DKT consistently outperforms the others across our four diverse datasets. However, we also find that our logistic regression model Combined-AugmentedLR, which combines the three above points, outperforms all of these neural network models on average across the four datasets, and outperforms them all on three of the four individual datasets (DKT outperforms Combined-AugmentedLR on the `Eedi` dataset). We do find that neural network approaches are promising, however, especially due to their ability in principle to automatically discover additional features that logistic regression cannot discover on its own. Furthermore, we believe this ability of neural networks will be increasingly important as available student log datasets continue to increase both in size and in diversity of logged features. We conclude that a promising direction for future research is to explore the integration of our above three approaches into DKT and other neural network approaches.

It is useful to consider how our results relate to previously published results. We found that the time window feature sets proposed by DAS3H (Choffin et al., 2019) enhanced Best-LR predictions for all four datasets, showing that these features can also improve other student performance models besides the DASH model considered in the original DAS3H paper. Note that when Gervet et al. 2020 introduced the Best-LR model they also experimented with time window features, but unlike us did not observe consistent benefits. This might be due to the number of additional parameters that must be trained when incorporating these time window features, and

the corresponding need for larger training datasets. Recall the datasets we used in this manuscript are about one order of magnitude larger than the ones used by [Gervet et al. 2020](#). Our algorithm comparison (Table 5) revealed that the prediction quality of PPE ([Walsh et al., 2018](#)) is worse than all other considered student performance modeling techniques. This is likely due to the fact that PPE was designed to model cognitive processes during word pair learning over longer periods of time. This is a very different setting than the learning experiences offered by the four studied tutoring systems. Three of them focus on mathematics and EdNet KT3 prepares students for the TOEIC<sup>®</sup> examination which goes beyond conventional retrieval practice. The elapsed and lag time features introduced by SAINT+ ([Shin et al., 2021](#)) also improve Best-LR predictions substantially in our experiments. Interestingly, the performance increment for the Best-LR model produced by these features is comparable to and sometimes even greater than the performance difference between SAINT and SAINT+. Note SAINT+ is an improved version of SAINT that uses the two interaction time based features. This suggests it might not require a deep learning-based approach to leverage these elapsed time and lag time features optimally.

Considering features that require augmented student logs that contain more than question-response pairs, we found these augmented features vary in utility. We found the feature "current study module" to be a particularly useful signal for all datasets. During preliminary analysis we observed differences in student performance between different parts of the learning sessions (i.e., pre-test, learning, post-test, . . .), which are captured by the study module feature. Even though post-tests tend to contain more difficult questions on average, the highest level of student performance is observed during post-test session in which the users' overall performance is evaluated.

Importantly, we also found that introducing background knowledge about prerequisites and post-requisites among the knowledge components (KCs) in the curriculum is very useful. As summarized in Table 3, counts of correctly answered questions and attempts associated with pre- and post-requisite KCs are among the most informative features. Importantly, these features can be easily incorporated even into pre-existing ITSs and datasets that log only question-response pairs, because calculating these features requires no new log data – only annotating it using background knowledge about their prerequisite structure.

We compared different types of machine learning algorithms for student performance modeling in Subsection 5.3. One limitation is that we did not refit IRT's student ability parameter after each user response. [Wilson et al. 2016](#) showed that refitting the ability parameter after each interaction makes IRT more competitive on multiple smaller datasets. While our feature selection for the AugmentedLR models solely focused on achieving accurate performance predictions, the inclusion of certain contextual features can lead to reductions in generalizability to unseen data. For example, the use of school or teacher specific parameters requires us to refit the model periodically as new schools and teachers start working with the system. When selecting a feature set for real world applications one might want to trade a small reduction in predictive performance in favor of enhanced generalizability to new users. While the technique of using multiple time-specialized performance models can improve the overall prediction accuracy when using the Best-LR and Best-LR+ feature sets, the approach can also harm the predictive performance of some of the models that use the higher-parametric AugmentedLR feature set. This is likely due to the overfitting problem associated with more expressive modeling approaches. The proposed technique represents a first step towards mitigating the performance modeling cold-start problem for new students and we see a lot of potential for future research in this direction.

Recently, multiple intricate deep learning based techniques have been proposed and yield

state-of-the-art performance for specific datasets (e.g., [Shin et al. 2021](#), [Zhou et al. 2021](#), [C. Zhang et al. 2021](#), ...). Unfortunately, many of these works only employ one or two datasets, which raises the question of how suitable they are for other tutoring systems. The code and new data released alongside this manuscript increases the usability of multiple large-scale educational datasets. We hope that future works will leverage these available datasets to test whether novel student performance modeling algorithms are effective across different tutoring systems.

## 7 CONCLUSION

In this paper we present several approaches to extending the state of the art in student performance modeling. We show that approaches based on logistic regression can be improved by adding specific new features to student log data, including features that can be calculated from simple question-response student logs, and features not traditionally captured in student logs such as which videos the student watched or skipped. Furthermore, we introduce a new approach of training multiple logistic regression models on different partitions of the data, and show further improvements in student performance modeling by automatically combining the predictions of these multiple specialized logistic regression models.

Taken together, our proposed improvements lead to our Combined-AugmentedLR method which achieves a new state of the art for student performance modeling. Whereas the previous state-of-the-art logistic regression approach, Best-LR ([Gervet et al., 2020](#)), achieves an AUC score of 0.767 on average over our four datasets, our Combined-AugmentedLR approach achieves an improved AUC of 0.808 – a reduction of 17.5% in the AUC error (i.e., in the difference between the achieved AUC score and the ideal AUC score of 1.0). Furthermore, we observed that Combined-AugmentedLR achieves improvements consistently across all four of the diverse datasets we considered, suggesting that our methods should be useful across a broad range of intelligent tutoring systems.

To encourage researchers to compare against our approach, three of the four datasets we chose are publicly available. In addition, we make the entire code base for the algorithms and experiments reported here available on GitHub<sup>3</sup>. Our implementation converts the four large-scale datasets into a standardized format and uses parallelization for efficient processing. It increases the usability of the large-scale log data, and we hope that it will benefit future research on novel student performance modeling techniques.

## A APPENDIX: EXTENDED FEATURE DESCRIPTION

This Appendix provides a reference with additional implementation details for the features described in Section 4. The individual ITSs capture similar information in different ways and we discuss necessary system specific adaptations. While we are already sharing the complete code base that was used to generate the experimental results in this paper on GitHub, this Appendix is intended to guide independent re-implementations.

### QUESTION (ITEM) ID

All four of our datasets assign each question (i.e., item) a unique identifier. To make a performance prediction for a question our implementation converts the corresponding question identifier into a

---

<sup>3</sup><https://github.com/rschmucker/Large-Scale-Knowledge-Tracing>

sparse one-hot vector which is then passed to the machine learning algorithm. Note a one-hot vector refers to a vector where all values are zero except for a single value of 1. For example, given a dataset with  $n$  different questions, our one-hot vector representation contains  $n$  values, of which  $n - 1$  are zeros, and just a single value of 1 to indicate the current question. Knowledge about question identifiers allows a model to learn question specific difficulty parameters.

## KNOWLEDGE COMPONENT (KC) ID

Knowledge about KC identifiers allows a model to learn skill specific difficulty parameters. While `ElemMath2021` and `Junyi15` only assign a single KC identifier per question, `Ednet`, `KT3` and `Eedi` can assign multiple KC identifiers to the same question. To make a performance prediction for a question our implementation uses a sparse-vector which is 0 everywhere except in the entries which mark the corresponding KC identifiers with a 1 value.

## TOTAL/KC/QUESTION COUNTS

Count features summarize a student's history of past interaction with the ITS and are an important component of PFA (Pavlik Jr et al., 2009) and Best-LR (Gervet et al., 2020). In our experiments we evaluate three ways of counting the number of prior correct responses and overall attempts:

1. Total counts: Here we compute two features capturing the total number of prior correct responses and overall attempts.
2. KC counts: For each individual KC we compute two features capturing the number of prior correct responses and attempts related to questions that target the respective KC. A vector containing the counts for all KCs related to the current question is then passed to the machine learning algorithm.
3. Question counts: Here we compute two features capturing the total number of prior correct responses and attempts on the current question.

All count features are subjected to scaling function  $\phi(x) = \log(1 + x)$  before being passed to the machine learning algorithm. This avoids features of large magnitude.

## TOTAL/KC/QUESTION TIME-WINDOW (TW) COUNTS

Time-window based count features summarize student history over different periods of time and provide the model with temporal information (Choffin et al., 2019). Following the original DAS3H implementation we define a set of time windows  $W = \{1/24, 1, 7, 30, +\infty\}$  measured in days. For each window  $w \in W$ , we count the number of prior correct responses and overall attempts of the student which fall into the window. We evaluate three ways of counting the number of prior correct responses and overall attempts:

1. Total TW counts: For each time-window, we compute two features capturing the total number of prior correct responses and overall attempts.
2. KC TW counts: For each time-window and each individual KC, we compute two features capturing the number of prior correct responses and attempts related to questions that target the respective KC. A vector containing the counts for all time-windows and KCs related to the current question is then passed to the machine learning algorithm.



3. Question TW counts: For each time-window, we compute two features capturing the total number of prior correct responses and attempts on the current question.

All count features are subjected to scaling function  $\phi(x) = \log(1 + x)$  before being passed to the machine learning algorithm. This avoids features of large magnitude.

### R-PFA $F$ & $R$

Motivated by the idea that more recently observed student responses are more indicative for future performance than older ones, R-PFA (Galyardt and Goldin, 2015) augments PFA (Pavlik Jr et al., 2009) by introducing two new features. For each KC  $k$  R-PFA considers all interactions of student  $s$  with  $k$  up to time  $t$  and computes: (i) A recency-weighted count of previous failures  $F_{s,k,t}$  using exponential decay. (ii) A recency-weighted proportion of past successes  $R_{s,k,t}$  using normalized exponential decay. The degree of decay is controlled by the hyperparameters  $d_F$  and  $d_R \in [0, 1]$ . To allow the computation of  $R_{s,k,t}$  when a student visits a KC  $k$  for the first time, their interaction history is appended with  $g = 3$  incorrect “ghost attempt”. The total number of responses of student  $s$  related to KC  $k$  is  $a_{s,k}$  and correctness indicator  $a_{s,k,i}$  is 1 when  $s$ 's  $i$ -th attempt on KC  $k$  was correct and 0 otherwise. With this the two R-PFA features are defined as

$$F_{s,k,t} = \sum_{i=1}^{a_{s,k}} d_F^{(a_{s,k}+1)-i} (1 - a_{s,k,i}), \quad R_{s,k,t} = \sum_{i=(1-g)}^{a_{s,k}} \frac{d_R^{a_{s,k}-i}}{\sum_{j=(1-g)}^{a_{s,k}} d_R^{a_{s,k}-j}} a_{s,k,i}. \quad (9)$$

### PPE COUNT

In the context of vocabulary learning, PPE (Walsh et al., 2018) was proposed to capture the spacing effect (Cepeda et al., 2008). PPE does so by introducing a weighting scheme which considers the delay between individual practice session and by assuming a multiplicative relationship between the number of prior attempts  $a_{s,k}$  with a time variable  $T_k$ . Further, the model uses a learning rate parameter  $c$  and three forgetting rate parameters  $x$ ,  $b$  and  $m$ . These four hyperparameters need to be set by the user. Let  $\Delta_{s,k,i}$  be the real time passed since student  $s$ 's  $i$ -th response to KC  $k$ . For our feature evaluation define a weighted count feature  $B_{s,k} = a_{s,k}^c T_k^{-d_t}$  by using PPE's multiplicative relationship between  $a_{s,k}$  and  $T_k$  which are defined as

$$T_k = \left( \sum_{i=1}^{a_{s,k}} \Delta_{s,k,i}^{1-x} \right) \left( \sum_{i=1}^{a_{s,k}} \frac{1}{\Delta_{s,k,i}^{-x}} \right), \quad d_t = b + m \left( \frac{1}{a_{s,k}} \sum_{i=1}^{a_{s,k}} \frac{1}{\ln(\Delta_{s,k,i} - \Delta_{s,k,i+1} + e)} \right). \quad (10)$$

### CURRENT/PRIOR ELAPSED TIME

Elapsed time measures the time span from question display to response submission (Shin et al., 2021). The idea is that a faster response is correlated with student proficiency. For our experiments we subject elapsed time values to scaling function  $\phi(x) = \log(1 + x)$  and also use the categorical one-hot encodings from (Shin et al., 2021). There, elapsed time is capped off at 300 seconds and categorized based on the integer second. We evaluate two variations of this feature. In the first version we compute elapsed time based on interactions with the *current* question. In the second version we compute elapsed time based on interactions with the *prior* question. Because it is unknown how long a student will take to answer a question ahead of time the elapsed time value of the current question is not available for question scheduling purposes and is excluded from the model comparison in Subsection 5.3.



## CURRENT/PRIOR LAG TIME

Lag time measures the time passed between the completion of the previous exercise until the next question is received (Shin et al., 2021). Lag time can be indicative for short-term memorization and forgetting. For our experiments we subject lag time values to scaling function  $\phi(x) = \log(1 + x)$  and also use the categorical one-hot encodings from (Shin et al., 2021). There, lag time is rounded to integer minutes and assigned to one of 150 categories (0, 1, 2, 3, 4, 5, 10, 20, 30, . . . , 1440). Because we cannot compute a lag time value for the very first question a student encounters we use an additional indicator flag. We evaluate two variations of this feature. In the first version we compute lag time based on interactions with the *current* question. In the second version we compute lag time based on interactions with the *prior* question.

## DATE-TIME: MONTH, WEEK, DAY, HOUR

Date-time features provide information related to the temporal context a learning activity is placed in. Here we consider the month, week, day and hour of interaction. We pass each of these four attributes to the algorithm using one-hot encodings.

## STUDY MODULE: ONE-HOT/COUNTS

All four datasets group study activities into distinct categories or modules (e.g., pre-test, effective learning, review, . . .). Providing a machine learning model with information about the corresponding study module can help with adapting the predictions to the different learning contexts. ElemMath2021 indicates different modules with the *s\_module* attribute. EdNet KT3 indicates different modules with the *source* attribute. Eedi indicates different modules with the *SchemeOfWorkId* attribute. For Junyi15 we can derive 8 study module identifiers corresponding to the unique combinations of *topic\_mode*, *review\_mode* and *suggested* flags. We encode these study module attribute values into one-hot vectors before passing them to the algorithm.

## TEACHER/GROUP ID

The ElemMath2021 dataset annotates each student response with the identifier of the supervising teacher. Similarly, the Eedi dataset annotates each student response with their group identifier. Both attributes provide information about the current learning context. We encode the identifiers into one-hot to allow the machine learning algorithm to learn teacher/group specific parameters.

## SCHOOL ID

The ElemMath2021 dataset associates each student with a physical or virtual tutoring center. Each tutoring center is assigned a unique school identifier. To capture potential differences between the various schools we allow the machine learning algorithm to learn school specific parameters by encoding the school identifiers into one-hot.

## COURSE

The ElemMath2021 dataset contains logs from a variety of different mathematics courses. While each course is assigned a unique course identifier, the KCs and questions treated in the

individual courses can overlap. To capture differences in the context set by the individual courses we allow the machine learning algorithm to learn course specific parameters by encoding the course identifiers into one-hot.

### TOPIC

The `ElemMath2021` dataset organizes learning activities into courses which itself are split into multiple topics. Each topic is assigned a unique topic identifier. The ITS can deploy the same question in multiple topics and the differences in learning context might affect student performance. We allow the machine learning algorithm to learn topic specific parameters by encoding the topic identifiers into one-hot.

### DIFFICULTY

The `ElemMath2021` dataset associates each question with a manually assigned difficulty score from the set  $\{10, 20, 30, 40, 50, 60, 70, 80, 90\}$ . We learn a model parameter for each distinct difficulty value by using a one-hot encoding.

### BUNDLE/QUIZ ID

The `EdNet_KT3` datasets annotates each response with a bundle identifier and the `Eedi` dataset annotates each response with a quiz identifier. Both bundles and quizzes mark sets of multiple questions which are asked together. The ITS can decide to assign a bundle/quiz to the student which then needs to respond to all associated questions. To capture the learning context provided by the current bundle/quiz we encode the corresponding identifiers into one-hot.

### PART/AREA ONE-HOT/COUNTS

The `Ednet_KT3` dataset assigns each question one label based on which of the seven TOEIC<sup>®</sup> exam parts it addresses. Similarly, the `Junyi15` dataset assigns each questions one of nine area identifiers which marks the area of mathematics the question addresses. We allow the machine learning algorithm to learn part/area specific parameters by encoding the part/area identifiers into one-hot. In addition, we experiment with two count features capturing the total number of prior correct responses and attempts on questions related to the current part/area. Before passing the count features to the algorithm we subject them to scaling function  $\phi(x) = \log(1 + x)$ .

### AGE

The `Eedi` dataset provides an attribute which captures students' birth date. We learn a model parameter for each distinct age by using a one-hot encoding. Students without a specified age are assigned a separate parameter.

### GENDER

The `Eedi` dataset categorizes student gender into female, male, other, and unspecified. We learn a model parameter for each attribute value by using one-hot vectors of dimension four.

## SOCIAL SUPPORT

The `Eedi` dataset provides information on whether students qualify for England’s pupil premium grant (a social support program for disadvantaged students). The attribute categorizes students into qualified, unqualified, and unspecified. We learn a model parameter for each attribute value by using one-hot vectors of dimension three.

## PLATFORM

The `ElemMath2021` dataset contains an attribute which indicates if a question was answered from a physical tutoring center or the online system. Similarly, the `EdNet_KT3` dataset indicates if a question was answered using the mobile app or a web browser. To pass this information to the model we use a two-dimensional one-hot encoding.

## PREREQUISITE: ONE-HOT/COUNTS

In addition to a KC model three of the datasets provide a graph structure that captures semantic dependencies between individual KCs and questions. `ElemMath2021` offers a prerequisite graph that marks relationships between KCs. `Junyi15` provides a prerequisite graph that describes dependencies between questions. In contrast, `Eedi` organizes its KCs via a 4-level topic ontology tree. For example the KC *Add and Subtract Vectors* falls under the umbrella of *Basic Vectors* which itself is assigned to *Geometry and Measure* which is connected to the tree root *Mathematics*. To extract prerequisite features from `Eedi`’s KC ontology we derive a pseudo prerequisite graph by first taking the two lower layers of the ontology tree and then using the parent nodes as prerequisites to the leaf nodes. We evaluate two ways of utilizing prerequisite information for student performance modeling:

1. Prerequisite IDs: For each question we employ a sparse vector that is zero everywhere except in the entries that mark the relevant prerequisite KCs (for `ElemMath2021` and `Eedi`) or questions (for `Junyi15`).
2. Prerequisite counts: For each question we look at its prerequisite KCs (for `ElemMath2021` and `Eedi`) or prerequisite questions (for `Junyi15`). For each prerequisite we then compute two features capturing the number of prior correct responses and attempts related to the respective prerequisite KC or question. After being subjected to scaling function  $\phi(x) = \log(1+x)$ , a vector containing the counts for all relevant prerequisite KCs/questions is passed to the machine learning algorithm.

## POST-REQUISITE: ONE-HOT/COUNTS

By inverting the directed edges of the prerequisite graph, we derive a post-requisite graph. Analogous to the prerequisite case we encode post-requisite information in two ways: (i) As sparse vectors that are zero everywhere except in the entries that mark the post-requisite KCs/questions with a 1; (ii) As correct and attempt count features computed for each KC/question that is post-requisite to the current question. For further details refer to the above prerequisite feature description.

## VIDEO: COUNT/SKIPPED/TIME

The ElemMath2021 and EdNet KT3 dataset both provide information on how students interact with lecture videos. We evaluate three ways of utilizing video consumption behavior for performance modeling:

1. Videos watched count: Here we compute two features: (i) The total number of videos a student has interacted with before; (ii) The number of videos a student has interacted with before related to the KCs of the current question.
2. Videos skipped count: ElemMath2021 captures video skipping events directly. For EdNet KT3 we count a video as skipped if the student watches less than 90%. Again we compute two features: (i) The total number of videos a student has skipped before; (ii) The number of videos a student has skipped before related to the KCs of the current question.
3. Videos watched time: Here we compute two features: (i) The total time a student has spent watching videos in minutes; (ii) The time a student has spent watching videos related to the KCs of the current question in minutes.

All count and time features are subjected to scaling function  $\phi(x) = \log(1 + x)$  before being passed to the machine learning algorithm. This avoids features of large magnitude.

## READING: COUNT/TIME

The ElemMath2021 and EdNet KT3 dataset both provide information on how users interact with reading materials. ElemMath2021 captures when a student goes through a question analysis and EdNet KT3 creates a log whenever a student enters a written explanation. We evaluate two ways of utilizing reading behavior for student performance modeling:

1. Reading count: Here we compute two features: (i) The total number of reading materials a student has interacted with before; (ii) The number of reading materials a student has interacted with before related to the KCs of the current question.
2. Reading time: Here we compute two features: (i) The total time a student has spent on reading materials in minutes; (ii) The time a student has spent on reading materials related to the KCs of the current question in minutes.

The count and time features are subjected to scaling function  $\phi(x) = \log(1 + x)$  before being passed to the machine learning algorithm. This avoids features of large magnitude.

## HINT: COUNT/TIME

The Junyi15 dataset captures how students make use of hints. Whenever a student answers a question the system logs how many hints were used and how much time was spent on each individual hint. Students are allowed to submit multiple answers to the same question, though a correct response is only registered if it is the first attempt and no hints are used. We evaluate two ways of utilizing hint usage for student performance modeling:

1. Hint count: Here we compute two features: (i) The total number of hints a student has used before; (ii) The number of hints a student has used before related to the KCs of the current question.

2. Hint time: Here we compute two features: (i) The total time a student has spent on hints in minutes; (ii) The time a student has spent on hints related to the KCs of the current question in minutes.

The count and time features are subjected to scaling function  $\phi(x) = \log(1 + x)$  before being passed to the machine learning algorithm. This avoids features of large magnitude.

### SMOOTHED AVERAGE CORRECTNESS

Let  $c_s$  and  $a_s$  be the number of prior correct responses and overall attempts of student  $s$  respectively. Let  $\bar{r}$  be the average correctness rate over all other students in the dataset. The linear logistic model is unable to infer the ratio  $c_s/a_s$  of average student correctness on its own. Because of this we introduce the smoothed average correctness feature  $r_s$  to capture the average correctness of student  $s$  over time as

$$\tilde{r}_s = \frac{c_s + \eta\bar{r}}{a_s + \eta}.$$

Here,  $\eta \in \mathbb{N}$  is a smoothing parameter which biases the estimated average correctness rate,  $\tilde{r}_s$  of student  $s$  towards this all-students average  $\bar{r}$ . The use of smoothing reduces the feature variance during a student’s initial interactions with the ITS. A prior work by [Pavlik et al. 2021](#) proposed, but did not evaluate, an average correctness feature without the smoothing parameter for student performance modeling. While calibrating this parameter for our experiments, we observed benefits from smoothing and settled on  $\eta = 5$ .

### RESPONSE PATTERN

Inspired by the use of n-gram models in the NLP community (e.g., [Manning and Schütze 1999](#)), we propose *response patterns* as a feature which allows logistic regression models to infer factors impacting short-term student performance. At time  $t$ , a response pattern  $r_t \in \mathbb{R}^{2^n}$  is defined as a one-hot encoded vector that represents a student’s sequence of  $n \in \mathbb{N}$  most recent responses  $w_t = (a_{t-n}, \dots, a_{t-1})$  formed by binary correctness indicators  $a_{t-n}, \dots, a_{t-1} \in \{0, 1\}$ . The encoding process is visualized by Figure 3 in the paper.

## B APPENDIX: HYPERPARAMETERS AND MODEL SPECIFICATIONS

Here we provide information about the hyperparameter search spaces used for the deep learning based student performance modeling approaches we evaluated with a grid search in Subsection 5.3. For each model we defined a hyperparameter search space that captures and extends the hyperparameters which were used in the cited references (DKT, ([Piech et al., 2015](#)); SAKT, ([Pandey and Karypis, 2019](#)); SAINT, ([Choi et al., 2020](#)); SAINT+, ([Shin et al., 2021](#))). A detailed list of the used hyperparameter spaces is provided in Table 10. All models were trained for 100 epochs without learning rate decay.

### ACKNOWLEDGMENTS

We would like to thank Squirrel Ai Learning for providing the Squirrel Ai ElemMath2021 dataset. In particular, Richard Tong and Dan Bindman were instrumental in providing this dataset, in helping us understand it. We thank Jack Wang and Angus Lamb, for personal communication

Table 10: Hyperparameter search spaces for deep learning based approaches.

Model	DKT	SAKT	SAINT	SAINT+
Hidden & Embedding Size	{50, 100, 200, 500}	{50, 100, 200, 500}	{64, 128, 256, 512}	{64, 128, 256, 512}
Number of Layers	{1, 2}	{1, 2}	{2, 4, 6}	{2, 4, 6}
Dropout Rate	{0, 0.2, 0.5}	{0, 0.2, 0.5}	-	-
Truncated Sequence Length	-	200	100	100
Number of Heads	-	5	8	8
Learning Rate	$1 \times 10^{-3}$	$1 \times 10^{-3}$	$1 \times 10^{-3}$	$1 \times 10^{-3}$
Batch Size	{128, 256}	{128, 256}	{128, 256}	{128, 256}

enabling us to reconstruct the exact sequence of student interactions in the Eedi data logs. We thank Theophile Gervet for making his earlier implementations of several algorithms available in his code repository, and for help understanding the details. We are grateful to the teams at the CK-12 Foundation and at Squirrel Ai Learning for suggestions on how to improve the presentation in this paper. This work was supported in part through the CMU - Squirrel Ai Research Lab on Personalized Education at Scale, and in part by AFOSR under award FA95501710218.

## REFERENCES

- AIKENS, N. L. AND BARBARIN, O. 2008. Socioeconomic differences in reading trajectories: The contribution of family, neighborhood, and school contexts. *Journal of Educational Psychology* 100, 2, 235.
- BADRINATH, A., WANG, F., AND PARDOS, Z. 2021. pyBKT: An accessible python library of bayesian knowledge tracing models. In *Proceedings of the 14th International Conference on Educational Data Mining*. EDM, Paris, France, 468–474.
- BAKER, R., CORBETT, A. T., AND ALEVEN, V. 2008. More accurate student modeling through contextual estimation of slip and guess probabilities in bayesian knowledge tracing. In *Proceedings of the 9th International Conference on Intelligent Tutoring Systems*. Springer, Berlin, Germany, 406–415.
- BREIMAN, L. 2001. Random forests. *Machine learning* 45, 1, 5–32.
- CEN, H., KOEDINGER, K., AND JUNKER, B. 2006. Learning factors analysis—a general method for cognitive model evaluation and improvement. In *Proceedings of the 8th International Conference on Intelligent Tutoring Systems*. Springer, Berlin, Germany, 164–175.
- CEPEDA, N. J., VUL, E., ROHRER, D., WIXTED, J. T., AND PASHLER, H. 2008. Spacing effects in learning: A temporal ridgeline of optimal retention. *Psychological science* 19, 11, 1095–1102.
- CHANG, H.-S., HSU, H.-J., AND CHEN, K.-T. 2015. Modeling exercise relationships in e-learning: A unified approach. In *Proceedings of the 12th International Conference on Educational Data Mining*. EDM, Madrid, Spain, 532–535.
- CHOFFIN, B., POPINEAU, F., BOURDA, Y., AND VIE, J.-J. 2019. DAS3H: Modeling student learning and forgetting for optimally scheduling distributed practice of skills. In *Proceedings of the 12th International Conference on Educational Data Mining*. EDM, Montréal, QC, Canada, 29–38.
- CHOI, Y., LEE, Y., CHO, J., BAEK, J., KIM, B., CHA, Y., SHIN, D., BAE, C., AND HEO, J. 2020. Towards an appropriate query, key, and value computation for knowledge tracing. In *Proceedings of 7th ACM Conference on Learning @ Scale*. L@S ’20. ACM, New York, NY, USA, 341–344.
- CHOI, Y., LEE, Y., SHIN, D., CHO, J., PARK, S., LEE, S., BAEK, J., BAE, C., KIM, B., AND HEO, J. 2020. EdNet: A large-scale hierarchical dataset in education. In *Proceedings of the 21st*



- International Conference on Artificial Intelligence in Education*. Springer International Publishing, Cham, Switzerland, 69–73.
- CORBETT, A. T. AND ANDERSON, J. R. 1994. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction* 4, 4, 253–278.
- DING, X. AND LARSON, E. C. 2019. Why deep knowledge tracing has less depth than anticipated. In *Proceedings of the 12th International Conference on Educational Data Mining*. EDM, Montréal, QC, Canada, 282–287.
- DING, X. AND LARSON, E. C. 2021. On the interpretability of deep learning based models for knowledge tracing. *CoRR abs/2101.11335*.
- EGLINGTON, L. G. AND PAVLIK, JR, P. I. 2019. Predictiveness of prior failures is improved by incorporating trial duration. *Journal of Educational Data Mining* 11, 2 (9), 1–19.
- FENG, M., HEFFERNAN, N., AND KOEDINGER, K. 2009. Addressing the assessment challenge with an online system that tutors as it assesses. *User Modeling and User-Adapted Interaction* 19, 3, 243–266.
- GALYARDT, A. AND GOLDIN, I. 2015. Move your lamp post: Recent data reflects learner knowledge better than older data. *Journal of Educational Data Mining* 7, 2, 83–108.
- GERVET, T., KOEDINGER, K., SCHNEIDER, J., AND MITCHELL, T. 2020. When is deep learning the best approach to knowledge tracing? *Journal of Educational Data Mining* 12, 3, 31–54.
- GHOSH, A., HEFFERNAN, N., AND LAN, A. S. 2020. Context-aware attentive knowledge tracing. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD '20. ACM, New York, NY, USA, 2330–2339.
- GONZÁLEZ-BRENES, J., HUANG, Y., AND BRUSILOVSKY, P. 2014. General features in knowledge tracing: Applications to multiple subskills, temporal item response theory, and expert knowledge. In *Proceedings of 7th International Conference on Educational Data Mining*. EDM, London, UK, 84–91.
- HOCHREITER, S. AND SCHMIDHUBER, J. 1997. Long short-term memory. *Neural computation* 9, 8, 1735–1780.
- HUANG, X., CRAIG, S. D., XIE, J., GRAESSER, A., AND HU, X. 2016. Intelligent tutoring systems work as a math gap reducer in 6th grade after-school program. *Learning and Individual Differences* 47, 258–265.
- KHAJAH, M., LINDSEY, R. V., AND MOZER, M. C. 2016. How deep is knowledge tracing? In *Proceedings of the 9th International Conference on Educational Data Mining*. EDM, Raleigh, NC, USA, 94–101.
- KOEDINGER, K., SCHEINES, R., AND SCHALDENBRAND, P. 2018. Is the doer effect robust across multiple data sets? In *Proceedings of the 11th International Conference on Educational Data Mining*. EDM, Buffalo, NY, USA, 369–375.
- KOEDINGER, K. R., KIM, J., JIA, J. Z., MCLAUGHLIN, E. A., AND BIER, N. L. 2015. Learning is not a spectator sport: Doing is better than watching for learning from a mooc. In *Proceedings of the 2nd ACM Conference on Learning @ Scale*. L@S '15. ACM, New York, NY, USA, 111–120.
- KOEDINGER, K. R., MCLAUGHLIN, E. A., JIA, J. Z., AND BIER, N. L. 2016. Is the doer effect a causal relationship? how can we tell and why it's important. In *Proceedings of the 6th International Conference on Learning Analytics & Knowledge*. LAK '16. ACM, New York, NY, USA, 388–397.
- KÄSER, T., KLINGLER, S., SCHWING, A. G., AND GROSS, M. 2017. Dynamic bayesian networks for student modeling. *IEEE Transactions on Learning Technologies* 10, 4, 450–462.
- LECUN, Y., HAFFNER, P., BOTTOU, L., AND BENGIO, Y. 1999. *Object Recognition with Gradient-Based Learning*. Springer, Berlin, Germany, 319–345.

- LINDSEY, R. V., SHROYER, J. D., PASHLER, H., AND MOZER, M. C. 2014. Improving students' long-term knowledge retention through personalized review. *Psychological Science* 25, 3, 639–647.
- LIU, Q., HUANG, Z., YIN, Y., CHEN, E., XIONG, H., SU, Y., AND HU, G. 2019. Ekt: Exercise-aware knowledge tracing for student performance prediction. *IEEE Transactions on Knowledge and Data Engineering* 33, 1, 100–115.
- LIU, Q., SHEN, S., HUANG, Z., CHEN, E., AND ZHENG, Y. 2021. A survey of knowledge tracing. *CoRR abs/2105.15106*.
- MANNING, C. D. AND SCHÜTZE, H. 1999. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA, USA.
- MONTERO, S., ARORA, A., KELLY, S., MILNE, B., AND MOZER, M. 2018. Does deep knowledge tracing model interactions among skills? In *Proceedings of the 11th International Conference on Educational Data Mining*. EDM, Buffalo, NY, USA, 462–466.
- MOZER, M. C. AND LINDSEY, R. V. 2016. *Predicting and Improving Memory Retention: Psychological Theory Matters in the Big Data Era*. Routledge/Taylor & Francis Group, London, UK.
- NAKAGAWA, H., IWASAWA, Y., AND MATSUO, Y. 2019. Graph-based knowledge tracing: Modeling student proficiency using graph neural network. In *IEEE/WIC/ACM International Conference on Web Intelligence*. IEEE, New York, NY, USA, 156–163.
- PANDEY, S. AND KARYPIS, G. 2019. A self-attentive model for knowledge tracing. In *Proceedings of the 12th International Conference on Educational Data Mining*. EDM, Montréal, QC, Canada, 384–389.
- PARDOS, Z. A. AND HEFFERNAN, N. T. 2010. Modeling individualization in a bayesian networks implementation of knowledge tracing. In *Proceedings of the 18th International Conference on User Modeling, Adaptation, and Personalization*. Springer, Berlin, Germany, 255–266.
- PARDOS, Z. A. AND HEFFERNAN, N. T. 2011. KT-IDEM: Introducing item difficulty to the knowledge tracing model. In *Proceedings of the 19th International Conference on User Modeling, Adaptation, and Personalization*. Springer, Berlin, Germany, 243–254.
- PASZKE, A., GROSS, S., MASSA, F., LERER, A., BRADBURY, J., CHANAN, G., KILLEEN, T., LIN, Z., GIMELSHEIN, N., ANTIGA, L., ET AL. 2019. PyTorch: An imperative style, high-performance deep learning library. In *Proceedings of the 32th International Conference on Advances in Neural Information Processing Systems*. Vol. 32. Curran Associates, Inc., Vancouver, BC, Canada, 8026–8037.
- PAVLIK, P. I., EGLINGTON, L. G., AND HARRELL-WILLIAMS, L. M. 2021. Logistic knowledge tracing: A constrained framework for learner modeling. *IEEE Transactions on Learning Technologies* 14, 5, 624–639.
- PAVLIK JR, P., CEN, H., AND KOEDINGER, K. 2009. Performance factors analysis—a new alternative to knowledge tracing. In *Frontiers in Artificial Intelligence and Applications*. Vol. 200. IOS Press, Amsterdam, Netherlands, 531–538.
- PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A., MICHEL, V., THIRION, B., GRISEL, O., BLONDEL, M., PRETTENHOFER, P., WEISS, R., DUBOURG, V., ET AL. 2011. Scikit-learn: Machine learning in python. *The Journal of Machine Learning Research* 12, 2825–2830.
- PIECH, C., BASSEN, J., HUANG, J., GANGULI, S., SAHAMI, M., GUIBAS, L. J., AND SOHL-DICKSTEIN, J. 2015. Deep knowledge tracing. In *Proceedings of the 28th Conference on Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds. Vol. 28th. Curran Associates, Inc., Montréal, QC, Canada.
- POJEN, C., MINGEN, H., AND TZUYANG, T. 2020. Junyi academy online learning activity dataset: A large-scale public online learning activity dataset from elementary to senior high school students.

- [Available online at: <https://www.kaggle.com/datasets/junyiacademy/learning-activity-public-dataset-by-junyi-academy>; accessed 28-March-2022].
- QIU, Y., QI, Y., LU, H., PARDOS, Z. A., AND HEFFERNAN, N. T. 2011. Does time matter? modeling the effect of time with bayesian knowledge tracing. In *Proceedings of the 4th International Conference on Educational Data Mining*. Eindhoven University of Technology, Eindhoven, Netherlands, 139–148.
- RASCH, G. 1960. *Probabilistic Models for Some Intelligence and Attainment Tests*. Studies in mathematical psychology. Danmarks Paedagogiske Institut, Aarhus, Denmark.
- SAO PEDRO, M., BAKER, R., AND GOBERT, J. 2013. Incorporating scaffolding and tutor context into bayesian knowledge tracing to predict inquiry skill acquisition. In *Proceedings of the 6th International Conference on Educational Data Mining*. EDM, Memphis, TN, USA.
- SCRUGGS, R., BAKER, R., AND MCLAREN, B. 2020. Extending deep knowledge tracing: Inferring interpretable knowledge and predicting post system performance. In *Proceedings of the 28th International Conference on Computers in Education*. APSCE, Jhongli City, Taiwan, 195–204.
- SHEN, S., LIU, Q., CHEN, E., WU, H., HUANG, Z., ZHAO, W., SU, Y., MA, H., AND WANG, S. 2020. Convolutional knowledge tracing: Modeling individualization in student learning process. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, New York, NY, USA, 1857–1860.
- SHIN, D., SHIM, Y., YU, H., LEE, S., KIM, B., AND CHOI, Y. 2021. Saint+: Integrating temporal features for EdNet correctness prediction. In *Proceedings of the 11th International Learning Analytics and Knowledge Conference*. ACM, New York, NY, USA, 490–496.
- STAMPER, J., NICULESCU-MIZIL, A., RITTER, S., GORDON, G., AND KOEDINGER, K. 2010. Bridge to Algebra 2006-2007. development data set from KDD Cup 2010 educational data mining challenge. [Available online at: <https://pslcdatashop.web.cmu.edu/KDDCup/downloads.jsp>; accessed 28-March-2022].
- TONG, H., WANG, Z., LIU, Q., ZHOU, Y., AND HAN, W. 2020. HGKT : Introducing problem schema with hierarchical exercise graph for knowledge tracing. *CoRR abs/2006.16915*.
- TSUTSUMI, E., KINOSHITA, R., AND UENO, M. 2021. Deep-IRT with independent student and item networks. In *Proceedings of the 14th International Conference on Educational Data Mining*. EDM, Paris, France, 510–517.
- VAN CAMPENHOUT, R., JOHNSON, B., AND OLSEN, J. 2021. The doer effect: Replicating findings that doing causes learning. In *Proceedings of the 13th International Conference on Mobile, Hybrid, and On-line Learning*. IARIA, France, 1–6.
- VAN DER LINDEN, W. J. AND HAMBLETON, R. K. 2013. *Handbook of Modern Item Response Theory*. Springer, New York, NY, USA.
- VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L., GOMEZ, A. N., KAISER, L. U., AND POLOSUKHIN, I. 2017. Attention is all you need. In *Proceedings of the 31th Conference on Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. Vol. 30. Curran Associates, Inc., Long Beach, CA, USA.
- VON STUMM, S. 2017. Socioeconomic status amplifies the achievement gap throughout compulsory education independent of intelligence. *Intelligence* 60, 57–62.
- WALSH, M. M., GLUCK, K. A., GUNZELMANN, G., JASTRZEMBSKI, T., KRUSMARK, M., MYUNG, J. I., PITT, M. A., AND ZHOU, R. 2018. Mechanisms underlying the spacing effect in learning: A comparison of three computational models. *Journal of Experimental Psychology* 147, 9, 1325.

- WANG, Z., LAMB, A., SAVELIEV, E., CAMERON, P., ZAYKOV, J., HERNANDEZ-LOBATO, J. M., TURNER, R. E., BARANIUK, R. G., BARTON, C., JONES, S. P., ET AL. 2021. Results and insights from diagnostic questions: The neurips 2020 education challenge. In *NeurIPS 2020 Competition and Demonstration Track*. PMLR, Virtual, 191–205.
- WANG, Z., LAMB, A., SAVELIEV, E., CAMERON, P., ZAYKOV, Y., HERNÁNDEZ-LOBATO, J. M., TURNER, R. E., BARANIUK, R. G., BARTON, C., AND JONES. 2020. Diagnostic questions: The NeurIPS 2020 education challenge. *CoRR abs/2007.12061*.
- WHITE, K. R. 1982. The relation between socioeconomic status and academic achievement. *Psychological Bulletin* 91, 3, 461.
- WILSON, K. H., XIONG, X., KHAJAH, M., LINDSEY, R. V., ZHAO, S., KARKLIN, Y., VAN INWEGEN, E. G., HAN, B., EKANADHAM, C., BECK, J. E., ET AL. 2016. Estimating student proficiency: Deep learning is not the panacea. In *Neural Information Processing Systems, Workshop on Machine Learning for Education*. NIPS, Barcelona Spain, 3–7.
- WOLPERT, D. H. 1992. Stacked generalization. *Neural Networks* 5, 2, 241–259.
- YANG, H. AND CHEUNG, L. P. 2018. Implicit heterogeneous features embedding in deep knowledge tracing. *Cognitive Computation* 10, 1, 3–14.
- YANG, Y., SHEN, J., QU, Y., LIU, Y., WANG, K., ZHU, Y., ZHANG, W., AND YU, Y. 2021. Gikt: A graph-based interaction model for knowledge tracing. In *Machine Learning and Knowledge Discovery in Databases*. Springer International Publishing, Cham, Switzerland, 299–315.
- YEUNG, C.-K. 2019. Deep-IRT: Make deep learning based knowledge tracing explainable using item response theory. *CoRR abs/1904.11738*.
- YUDELSON, M. V., KOEDINGER, K. R., AND GORDON, G. J. 2013. Individualized bayesian knowledge tracing models. In *Proceedings of the 16th International Conference on Artificial Intelligence in Education*. Springer, Berlin, Germany, 171–180.
- ZHANG, C., JIANG, Y., ZHANG, W., AND GU, C. 2021. Muse: Multi-scale temporal features evolution for knowledge tracing. *CoRR abs/2102.00228*.
- ZHANG, J., DAS, R., BAKER, R., AND SCRUGGS, R. 2021. Knowledge tracing models’ predictive performance when a student starts a skill. In *Proceedings of the 14th International Conference on Educational Data Mining*. EDM, Paris, France, 625–629.
- ZHANG, J., SHI, X., KING, I., AND YEUNG, D.-Y. 2017. Dynamic key-value memory networks for knowledge tracing. In *Proceedings of the 26th International Conference on World Wide Web. WWW ’17. IW3C2*, Geneva, Switzerland, 765–774.
- ZHANG, L., XIONG, X., ZHAO, S., BOTELHO, A., AND HEFFERNAN, N. T. 2017. Incorporating rich features into deep knowledge tracing. In *Proceedings of 7th ACM Conference on Learning @ Scale. L@S ’17*. ACM, New York, NY, USA, 169–172.
- ZHAO, S., WANG, C., AND SAHEBI, S. 2020. Modeling knowledge acquisition from multiple learning resource types. In *Proceedings of the 13th International Conference on Educational Data Mining*. EDM, Virtual, 313–324.
- ZHOU, Y., LI, X., CAO, Y., ZHAO, X., YE, Q., AND LV, J. 2021. Lana: Towards personalized deep knowledge tracing through distinguishable interactive sequences. In *Proceedings of the 14th International Conference on Educational Data Mining*. EDM, Paris, France, 602–608.