

Modeling Knowledge Acquisition from Multiple Learning Resource Types

Siqian Zhao*
Computer Science
University at Albany, SUNY
Albany, NY 12222 USA
szhao2@albany.edu

Chunpai Wang*
Computer Science
University at Albany, SUNY
Albany, NY USA
cwang25@albany.edu

Shaghayegh Sahebi
Computer Science
University at Albany, SUNY
Albany, NY 12222 USA
ssahebi@albany.edu

ABSTRACT

Students acquire knowledge as they interact with a variety of learning materials, such as video lectures, problems, and discussions. Modeling student knowledge at each point during their learning period and understanding the contribution of each learning material to student knowledge are essential for detecting students' knowledge gaps and recommending learning materials to them. Current student knowledge modeling techniques mostly rely on one type of learning material, mainly problems, to model student knowledge growth. These approaches ignore the fact that students also learn from other types of material. In this paper, we propose a student knowledge model that can capture knowledge growth as a result of learning from a diverse set of learning resource types while unveiling the association between the learning materials of different types. Our multi-view knowledge model (MVKM) incorporates a flexible knowledge increase objective on top of a multi-view tensor factorization to capture occasional forgetting while representing student knowledge and learning material concepts in a lower-dimensional latent space. We evaluate our model in different experiments to show that it can accurately predict students' future performance, differentiate between knowledge gain in different student groups and concepts, and unveil hidden similarities across learning materials of different types.

Keywords

knowledge tracing, domain modeling, tensor factorization, multi-view learning

1. INTRODUCTION

Both student knowledge modeling and domain knowledge modeling are important problems in the educational data mining community. In this context, student knowledge tracing and knowledge modeling approaches aim to evaluate students' state of knowledge or quantify students' knowledge in

*First two authors contributed equally to this work.

the concepts that are presented in learning materials at each point of the learning period [15, 6, 51, 25, 53, 32, 14, 47]. Domain knowledge modeling, on the other hand, focuses on understanding and quantifying the topics, knowledge components, or concepts that are presented in the learning material [7, 12, 27]. It is useful in creating a coherent study plan for students, modeling students' knowledge, and analyzing students' knowledge gaps.

A successful student knowledge model should be personalized to capture individual differences in learning [51, 28], understand the association and relevance between learning from various concepts [42, 53], model knowledge gain as a gradual process resulting from student interactions with learning material [21, 38, 18], and allow for occasional forgetting of concepts in students [14, 32, 18]. Despite recent success in capturing these complexities in student knowledge modeling, a simple, but important aspect of student learning is still under-investigated: that students learn from different types of learning materials. Current research has focused on modeling one single type of learning resource at a time (typically, "problems"), ignoring the heterogeneity of learning resources from which students may learn. Modern online learning systems frequently offer students to learn and assess their knowledge using various learning resource types, such as readings, video lectures, assignments, quizzes, and discussions. Previous research has demonstrated considerable benefits of interacting with multiple types of materials on student learning. For example, worked examples can lead to faster and more effective learning compared to unsupported problem solving [33]; and enriching textbooks with additional forms of content, such as images and videos, increases the helpfulness of learning material [2, 1]. Ignoring diverse types of learning materials in student knowledge modeling limits our understanding of how students learn.

One of the obstacles in considering the combined effect of learning material types is the lack of explicit learning feedback from all of them. Some learning material types, such as problems and quizzes, are gradable. As students interact with such material types, the system can perceive student grade as an explicit feedback or indication of student knowledge: if a student receives a high grade in a problem, it is likely that the student has gained enough knowledge required to solve that problem. On the other hand, some of the learning materials are not gradable and their impact on student knowledge cannot be explicitly measured. For example, we cannot directly measure the consequent knowledge

Siqian Zhao, Chunpai Wang and Shaghayegh Sahebi "Modeling Knowledge Acquisition from Multiple Learning Resource Types" In: *Proceedings of The 13th International Conference on Educational Data Mining (EDM 2020)*, Anna N. Rafferty, Jacob Whitehill, Violetta Cavalli-Sforza, and Cristobal Romero (eds.) 2020, pp. 313 - 324

gain from watching a video lecture or studying an example.

As an alternative for quantifying student knowledge gain, the system can measure other quantities, such as the binary indication of student activity with a learning material or the time they spent on it. However, this kind of measure may result in contradictory conclusions [8, 23, 22]. For example, spending more time to study the examples provided by the system may both increase the student’s knowledge, and at the same time, be an indicator of a weaker student, who does not have enough knowledge in the provided concepts. These weaker students may select to study more examples to compensate for their lower knowledge levels. Consequently, the knowledge gain of studying these auxiliary learning materials is usually overpowered by the student selection bias and is not represented correctly in the overall dataset.

A similar issue exists in the current domain knowledge models. The automatic domain knowledge models that are based on students’ activities mainly model one type of learning material and ignore the relationship between various kinds of learning materials [17, 12]. Alternatively, an ideal domain knowledge model should be able to model and discover the similarities between learning materials of different types.

In this paper, we simultaneously address the problems of student knowledge modeling and domain knowledge modeling, while considering the heterogeneity of learning material types. We introduce a new student knowledge model that is the first to concurrently represent student interactions with both graded and non-graded learning material. Meanwhile, we discover the hidden concepts and similarities between different types of learning materials, as in a domain knowledge model. To do this, we pose this concurrent modeling as a multi-view tensor factorization problem, using one tensor for modeling student interactions with each learning material type. By experimenting on both synthetic and real-world datasets, we show that we can improve student performance prediction in graded learning materials, as measured by the Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE).

In summary, the contributions of this paper are:

- 1) proposing a personalized, multi-view student knowledge model (MVKM) that can capture learning from *multiple learning material types* and allow for occasional student forgetting, while modeling all types of learning materials;
- 2) conducting experiments on both synthetic and real-world datasets showing that our proposed model outperforms conventional methods in predicting student performance;
- 3) examining the resulting learning material and student knowledge latent features to show the captured similarity between learning material types and interpretability of student knowledge model.

2. RELATED WORK

Knowledge Modeling Student knowledge modeling aims to quantify student knowledge state in the concepts or skills that are covered by learning materials at each learning point.

Pioneer approaches of student knowledge modeling, despite being successful, were not personalized, relied on a predefined (sometimes expert-labeled) set of concepts in learning

material, did not allow for learned concepts to be forgotten by students, and modeled each concept independently from one another [19, 15, 36, 45]. Later, some student knowledge models aimed to solve these shortcomings by learning different parameters for each (type of) student [35, 51, 26], including decays to capture forgetting of concepts in learner models [39, 29, 31] and capturing the relationship between concepts that are present in a course [44, 21]. Yet, these models assume that a correct domain knowledge model, that maps learning material into course concepts, exists.

In recent years, new approaches aim to learn both domain knowledge model and student knowledge model at the same time [28, 20, 42, 48, 53, 18]. Our proposed model falls into this latest category as it does not require any manual labeling of learning materials, while having the ability to use such information if they are available. It is personalized by learning lower-dimensional student representations, allows forgetting of concepts during student learning by adding a rank-based constraint on student knowledge, and models the relationship between learning material.

Learning from Multiple Material Types In the educational data mining (EDM) literature, learning materials are provided in various types, such as problems, examples, videos, and readings. While there have been some studies in the literature on the value of having various types of learning materials for educating students [2, 8, 33], the relationship between these material types, and their combined effect on student knowledge and student performance is under-investigated.

Multiple learning material types have been studied in the literature in finding insights into different activity distributions or cluster patterns between high-grade and low-grade students [46, 49], have been used as contextual features in scaffolding or choosing among the existing student models [52, 43], have been added to improve existing domain knowledge models only for graded material types while ignoring student sequences [10, 13, 16, 30, 34, 41, 37], or have been classified into beneficial or non-beneficial for students [3]. However, to the best of our knowledge, none of these studies have explicitly modeled the contribution of various kinds of learning materials on student knowledge during the learning period, the interrelations among these learning materials, and their effect on student performance. The Bayesian Evaluation and Assessment framework found that assistance promoted students’ long-term learning. More recently, Huang et al. discovered that adaptation of their framework (FAST) for student modeling by including various activity types may lead researchers to contradictory conclusions [23]. More specifically, in one of their formulations student example activity suggests a positive association with model parameters, such as probability of learning, while in another formulation this type of activity has a negative association with model parameters. Also, Hosseini et al. concluded that annotated examples show a negative relationship with students’ learning, because of a selection effect: while annotated examples may help students to learn, weaker students may study more annotated examples [22]. The model proposed in this paper considers student interactions from multiple learning material types, mitigating over-estimation of student knowledge by transferring information from interactions with graded

material, while accounting for knowledge increase that happens as a result of student interaction with non-graded material.

3. MULTI-VIEW KNOWLEDGE MODELING

3.1 Problem Formulation and Assumptions

We consider an online learning system in which M students interact with and learn from multiple types ($r \in \mathfrak{R}$) of learning materials. Each learning material type r includes a set of $P^{[r]}$ learning materials. A material type can be either graded or non-graded. Students' normalized grade in tests, success or failure in compiling a piece of code, or scores in solving problems are all examples of graded learning feedback. Whereas, watching videos, posting comments in discussion forums, or interacting with annotated examples are instances of non-graded learning feedback that the system can receive. We model the learning period as a series of student attempts on learning materials, or time points ($a \in \mathcal{A}$). To represent student interaction feedback with learning materials of each type r during the whole learning period \mathcal{A} , we use a $M \times P^{[r]} \times A$ three-dimensional tensor $\mathbf{X}^{[r]}$. The a^{th} slice of tensor $\mathbf{X}^{[r]}$, denoted by $X_a^{[r]}$, is a matrix representing student interactions with the learning material type r during one snapshot of the learning period. The s^{th} row of this interaction matrix $\mathbf{x}_{a,s}^{[r]}$ shows feedback from student s 's interactions with all learning materials of type r at attempt a ; and the tensor element $x_{a,s,p}^{[r]}$ is the feedback value of student s 's activity on learning material p of type r at learning point a .

We use the following assumptions in our model: **(a)** Each learning material covers some concepts that are presented in a course; the set of all course concepts are shared across learning materials; and the training data does not include the learning materials' contents nor their concepts. **(b)** Different learning materials have different difficulty or helpfulness levels for students. For example, one quiz can be more difficult than another one, and one video lecture can be more helpful than the other one. **(c)** The course may follow a trend in presenting the learning material: going from easier concepts to more difficult ones or alternating between easy and difficult concepts; despite that, students can freely interact with the learning materials and are not bound to a specific sequence. **(d)** As students interact with these materials, they learn the concepts that are presented in them; meaning that their knowledge in these concepts increases. **(e)** Since students may forget some course concepts, this knowledge increase is not strict. **(f)** Different students come with different learning abilities and initial knowledge values. **(g)** The gradual change of knowledge varies among different students. But, students can be grouped together according to how their knowledge changes in different concepts, e.g., some students are fast learners compared to others. **(h)** Eventually, a student's performance in a graded learning material, represented by a score, depends on the concepts covered in that material, student's knowledge in those concepts, the learning material difficulty/helpfulness, and the general student ability.

In addition to the above, we have an essential assumption **(i)** that connects the different parts of our model: a student's knowledge that is obtained from interacting with one learning material type is transferable to be used in other types of

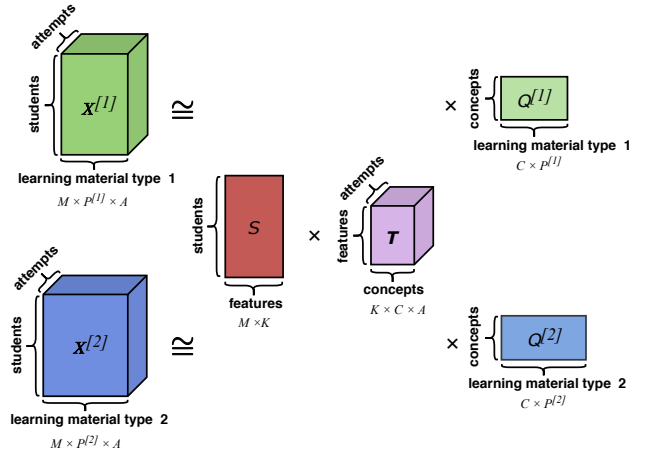


Figure 1: Decomposing student interaction tensors with two learning material types $\mathbf{X}^{[1]}$ and $\mathbf{X}^{[2]}$.

learning materials. In other words, students' knowledge can be modeled and quantified in the same latent space for all different learning material types. In the following, we first propose a single-view model for capturing the knowledge gained using one type of learning material (MVKM-Base) and then extend it to a multi-view model that can represent multiple types of learning materials.

3.2 MVKM Factorization Model

The Proposed Base Model (MVKM-Base). Following the mentioned assumptions in Section 3.1, particularly assumptions **(a)**, **(g)**, and **(h)**, and assuming that students interact with only one learning material type, we model student interaction tensor \mathbf{X} as a factorization (n -mode tensor product) of three lower-dimensional representations: 1) an $M \times K$ student latent feature matrix S , 2) a $K \times C \times A$ temporal dynamic knowledge tensor T , and 3) a $C \times P$ matrix Q serving as a mapping between learning materials and course concepts. In other words, we have $\hat{x}_{s,a,p} \approx \mathbf{s}_s \cdot T_a \cdot \mathbf{q}_p$. Matrix S here represents students being mapped to latent learning features that can be used to group the students (assumption **(g)**). Tensor T quantifies the knowledge growth of students with each learning feature in each of the concepts while attempting the learning material. Accordingly, the resulting tensor from product $\mathbf{K} = \mathbf{S}\mathbf{T}$ represents each student's knowledge in each concept at each attempt.

To increase interpretability, we enforce the contribution of different concepts in each learning material to be non-negative and sum to one. Similarly, we enforce the same constraints on each student's membership in the student latent features. Since each student can have a different ability (assumption **(f)**) and each learning material can have its own difficulty level (assumption **(b)**), we add two bias terms to our model (b_s for each student s , and b_p for each learning material p) to account for such differences. To capture the general score trends in the course (assumption **(c)**), we add a parameter b_a for each attempt. Accordingly, we estimate student s 's score in a graded learning material p at attempt a ($\hat{x}_{a,s,p}$) as in Equation 1. Here, T_a is a matrix capturing the relationship between student features and concepts at attempt a , \mathbf{s}_s represents student s 's latent feature vector, \mathbf{q}_p shows material p 's concept vector.

$$\hat{x}_{s,a,p} \approx \mathbf{s}_s \cdot T_a \cdot \mathbf{q}_p + b_s + b_p + b_a \quad (1)$$

We use a sigmoid function $\sigma(\cdot)$ to estimate student interaction with a non-graded learning material, or graded ones with binary feedback:

$$\hat{x}_{s,a,p} \approx \sigma(\mathbf{s}_s \cdot T_a \cdot \mathbf{q}_p + b_s + b_p + b_a)$$

Modeling Knowledge Gain while Allowing Forgetting. So far, this simple model captures latent feature vectors of students and learning materials, and learns T as a representation of knowledge in students. However, it does not explicitly model students’ gradual knowledge gain (assumption **(d)**). We note that students’ knowledge increase is associated with the strength of concepts in the learning material that they interact with. As students interact with learning materials with some specific concepts, it is more likely for their predicted scores in the relevant learning materials to increase. With a Markovian assumption, we can say that if students have practiced some concepts, we expect their scores in attempt $a + 1$ to be more than their scores in attempt a :

$$\mathbf{s}_s \cdot T_{a+1} \cdot \mathbf{q}_p - \mathbf{s}_s \cdot T_a \cdot \mathbf{q}_p \geq 0$$

However, this inequality constraint is too strict as the students may occasionally forget the learned concepts (assumption **(e)**). To allow for this occasional forgetting and soften this constraint, we model the knowledge increase as a rank-based constraint, that allows for knowledge loss, but penalizes it. We formulate this constraint as maximising the value for \mathcal{L}_2 in Equation 2. Essentially, this penalty term can be viewed as a prediction-consistent regularization. It helps to avoid significant changes in students’ knowledge level since their performance is expected to transit gradually over time.

$$\mathcal{L}_2 = \sum_{j=1}^{a-1} \sum_{s,p} \log(\sigma(\mathbf{s}_s \cdot T_a \cdot \mathbf{q}_p - \mathbf{s}_s \cdot T_j \cdot \mathbf{q}_p)) \quad (2)$$

The Proposed Multi-View Model (MVKM). We rely on our main assumption **(i)** to extend our model to capture learning from different learning material types. So far, we have assumed that course concepts are shared among learning materials (assumption **(a)**). With the knowledge transfer assumption **(i)**, all learning materials of different types will share the same latent space. Also, we represent student knowledge and student ability as shared parameters across all different learning material types. Consequently, for each set of learning materials of type $r \in \mathfrak{R}$, we can rewrite Equation 1 as:

$$\hat{x}_{s,a,p}^{[r]} \approx \mathbf{s}_s \cdot T_a \cdot \mathbf{q}_p^{[r]} + b_s + b_p^{[r]} + b_a$$

An illustration of this decomposition, when considering two learning material types, is presented in Figure 1. Note that we represent one shared matrix student S and one shared knowledge gain tensor T in both types of learning materials.

We can learn the parameters of our model by minimizing the sum of squared differences between the observed ($x_{s,a,p}^{[r]}$) and estimated ($\hat{x}_{s,a,p}^{[r]}$) values over all learning material types $r \in \mathfrak{R}$. For the learned parameters to be generalizable to unseen data, we regularize the unconstrained parameters using their L-2 norms. As a result, we minimize the objective function in Equation 3, in which $\gamma^{[r]}$ are hyper-parameters that represent the relative importance of different learning mate-

rials types. λ_t and λ_s are hyper-parameters to control the weights of regularization term of T and S .

$$\begin{aligned} \mathcal{L}_1 = \sum_{r,s,a,p} \gamma^{[r]} (\hat{x}_{s,a,p}^{[r]} - x_{s,a,p}^{[r]})^2 + \lambda_t \|T_a\|_F^2 + \lambda_s \|\mathbf{s}_s\|_F^2 \\ \text{s.t. } \forall r,c,p \quad q_{c,p}^{[r]} \geq 0, \quad \sum_c q_{c,p}^{[r]} = 1 \end{aligned} \quad (3)$$

Similarly, the knowledge gain and forgetting constraint presented in Equation 2 can be extended to the multi-view model. Eventually, we use a combination of the reconstruction objective function (Equation 3) and the learning and forgetting objective function (Equation 2) to model students’ knowledge increase, while representing their personalized knowledge and finding learning material latent features, as in Equation 4. Note that, since our goal is to minimize \mathcal{L}_1 and maximize \mathcal{L}_2 , we use $-\mathcal{L}_2$ to minimize \mathcal{L} . To balance between the accuracy of student performance prediction and modeling student knowledge increase, we use a nonnegative trade-off parameter ω :

$$\mathcal{L} = \mathcal{L}_1 - \omega \mathcal{L}_2 \quad (4)$$

We use stochastic gradient descent algorithm to minimize \mathcal{L} in Equation 4. The parameters need to learn are students’ latent feature matrix (S), dynamic knowledge in each concept at any attempt (T), importance of each concept in every learning material ($Q^{[r]}$), each student’s general ability (b_s), each learning material’s difficulty/helpfulness ($b_p^{[r]}$), and each attempt’s bias ($b_a^{[r]}$).

4. EXPERIMENTS

We evaluate our model with three sets of experiments. First, to validate if the model captures the variability of observed data, we use it to predict unobserved student performances (Sec. 4.3). Second, to check if our model represents valid student knowledge growth, we study the knowledge increase patterns between different types of students and across different concepts (Sec. 4.4). Finally, to study if the model meaningfully recovers learning materials’ latent concepts, we analyze their similarities according to the learned latent feature vectors (Sec. 4.5). Without loss of generalizability, although the model is designed to handle multiple learning material types, we experiment on two learning material types. Before the experiments, we will go over our datasets, and experiment setup.

Dataset	material type 1 (#)	material type 2 (#)	#stu	act. seq. len.	#rcds.	avg. sco.
Synthetic_NG	quiz (10)	discussion (15)	1000	20	19991	0.6230
Synthetic_NG2	quiz (10)	discussion (15)	1000	20	19991	0.6984
Synthetic_G	quiz (10)	assignment (15)	1000	20	19980	0.6255
MORF_QD	assignment (18)	discussion (525)	459	25	6800	0.8693
MORF_QL	assignment (10)	lecture (52)	1329	76	58956	0.7731
Canvas_H	quiz (10)	discussion (43)	1091	20	13633	0.8648

Table 1: Statistics for each datasets, where #stu is number of students, act. seq. len. is the maximum activity length, #rcds. is number of records that student interact with learning materials and avg. sco. is graded learning material’s average score.

4.1 Datasets

We use three synthetic and three real-world datasets (from two MOOCs) to evaluate the proposed model. Our choice of real-world datasets is guided by two factors, aligned with

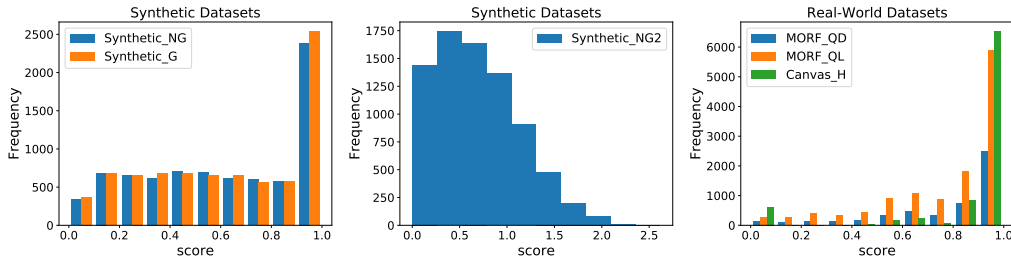


Figure 2: Histogram of graded materials' Scores in Synthetic Data and Real-World Data.

our assumptions: that they include multiple types of learning material, and that they allow the students to work freely with the learning material in the order they choose. In the real-world datasets, we select the students that work with both types of learning materials, removing the learning materials that none of these students have interacted with. General statistics of each dataset are presented in Table 1. Figure 2 shows score distributions of the graded learning material types in these datasets.

Synthetic Data. We generate three synthetic datasets according to two characteristics: (1) if both learning material types are graded vs. if one of them is non-graded (or has binary observations); (2) if the student scores are capped and their distribution is highly skewed vs. if the score distribution is not capped and less skewed.

For creating the datasets, we follow similar assumptions as to the ones made by our model. Expecting $P^{[1]}$ learning materials of type 1, and $P^{[2]}$ materials of type 2, we first generate a random sequence L_s for each student s , which represents the student's attempts on different learning materials. Considering C latent concepts, we then create two random matrices $Q^{[1]} \in \mathbb{R}^{C \times P^{[1]}}$ and $Q^{[2]} \in \mathbb{R}^{C \times P^{[2]}}$ as the mapping between the learning material and the C underlying concepts, such that the sum of values for each underlying learning material is one. For the student knowledge gain assumption, we represent each student's knowledge increase separately. Hence, we directly create a student knowledge tensor \mathbf{K} , instead of creating \mathbf{S} and \mathbf{T} , and multiplying them. To generate \mathbf{K} , we first generate a random matrix K_1 that represents all students' initial knowledge in all C concepts. For generating the knowledge matrix in the next attempts (K_a), we use the following random process. For each student s , we generate a random number α representing the probability of forgetting. If $\alpha > \theta$ (forgetting threshold), we assume no forgetting happens and increase the value in the knowledge matrix according to the learning material that the student has interacted with: $k_{s,a} = k_{s,a-1} + \beta q_{L_s[a]}^{[r]}$. Here, β is a random effect of increasing and $L_s[a]$ is the learning material that student has selected to interact with at timestamp a . Otherwise ($\alpha < \theta$, or forget), we set $k_{s,a,c} = k_{s,a-1,c} - \text{rand}(0, \epsilon)$, for $\forall c \in C$. we use n-mode tensor product to build $\mathbf{X}^{[1]}$ and $\mathbf{X}^{[2]}$, where $\mathbf{X}^{[1]} = \mathbf{K}Q^{[1]}$, $\mathbf{X}^{[2]} = \mathbf{K}Q^{[2]}$. Finally, according to the student learning sequences L_s , we remove the "unobserved" values that are not in L_s from $\mathbf{X}^{[1]}$ and $\mathbf{X}^{[2]}$.

To create different data types according to the first characteristic above, for the graded learning material type r , we keep the values in $\mathbf{X}^{[r]}$. For the non-graded ones, we use the same process as above, except that in the final step we

set $x_{s,a,p}^{[r]} = 1$ according to the student sequence L_s . However, in many real-world scenarios, the score distribution of students is highly skewed especially towards higher scores (Figure 2 show it). To represent this skewness, in some of the generated datasets, we clip all $x_{s,a,p}^{[r]} > 1$ to 1.

Then, we create following three datasets according to above process: *Synthetic_G*, in which both learning material types are graded and scores are skewed; *Synthetic_NG*, in which one of the learning material types is graded and scores are skewed; and *Synthetic_NG2*, in which one of the learning material types is graded and scores are not skewed. We generate all synthetic data with 1000 students, $P^{[1]} = 10$ learning materials of type 1, $P^{[2]} = 15$ learning materials of type 2, $C = 3$ latent concepts, and maximum sequence length of 20 for students.

Canvas Network [11]. This is an online available dataset collected from various courses on the Canvas network platform¹. The available open online course data comes from various study fields, such as computer science, business and management, and humanities. For each course, its general field of study is presented in the data. The rest of the dataset is anonymized such that course names, discussion contents, student IDs, submission contents, or course contents are not available. Each course can have different learning material types, including assignments, discussions, and quizzes. We experiment on the data from one course in this system, with course id 770000832960975, which is in the humanities field (Canvas_H dataset). We use quizzes as the graded learning material type and discussions as the non-graded one.

MORF [4]. This is a dataset of the "educational data mining" course [5] at Coursera², available via the MOOC Replication Framework (MORF). The course includes various learning material types, including video lectures, assignments, and discussion forums. Students' history, in terms of their watched video lectures, submitted assignments, and participated discussions, in addition to the score they received in assignments, is available in data. In this course, we experiment with two datasets, each focusing on two sets of learning material types: one with assignments as the graded type and discussions as the non-graded type (MORF_QD), another with assignments as the graded type and video lecture views as the non-graded type (MORF_QL).

4.2 Experiment Setup

We use 5-fold student-stratified cross-validation to separate our datasets into test and train. At each fold, we use interaction records from 80% of students as training data. For the

¹<http://canvas.net>

²<https://www.coursera.org/>

rest (20%) of the students (target students), we split their attempt sequences on the graded learning material type into two parts: the first 50% and the last 50%. For performance prediction experiments, we predict the performance of the graded learning material type in the last 50%, given the first 50%. In order to see how the proposed model captures the knowledge growth, we do online testing, in which we predict the test data attempt by attempt (next attempt prediction). Eventually, we report the average performance on all five folds. For selecting the best hyper-parameters, we use a separate validation dataset. Our code and synthetic data are available at [GitHub](#)³.

4.3 Student Performance Prediction

In this set of experiments, we test our model on predicting student scores on their future unobserved graded learning material attempts. More specifically, we estimate student scores on their future attempts, and compare them with their actual scores in the test data.

4.3.1 Baselines

We compare our model with state-of-the-art student performance prediction baselines:

Individualized Bayesian Knowledge Tracing (IBKT) [24, 51]: This is a variant of the standard BKT model, which assumes binary observations and provides individualization on student priors, learning rate, guess, and slip parameters⁴.

Deep Knowledge Tracing (DKT) [38]: DKT is a pioneer algorithm that uses recurrent neural networks to model student learning, on binary (success/failure) student scores.

Feedback-Driven Tensor Factorization (FDTF) [40]: This tensor factorization model decomposes the student interaction tensor into a learning material latent matrix and a knowledge tensor. However, it only models one type of learning material, does not capture student latent features, and does not allow the students to forget the learned concepts. It assumes that students' knowledge strictly increases as they interact with learning materials.

Tensor Factorization Without Learning (TFWL): This is a model similar as FDTF, the only difference is TFWL does not have constraint that student knowledge is increasing.

Rank-Based Tensor Factorization (RBTF) [18]: This model has similar assumptions to FDTF. Except, it allows for occasional forgetting of concepts and has extra bias terms. Compared to MVKM, it does not differentiate between different student groups. It only uses student previous scores in graded learning materials to predict students' future scores, and it has a different tensor factorization strategy.

Bayesian Probabilistic Tensor Factorization (BPTF) [50]: This is a recommender systems model has a smoothing assumption over student scores in consecutive attempts.

AVG: This baseline uses the average of all students' scores for all predictions.

As mentioned before, one major issue in real-world datasets is their skewness, meaning that, on average, student grades are skewed towards a full (complete) score on quizzes/assignments. This skewness adds to the complexity of predicting an accurate score for unobserved quizzes: only using an overall average score will provide a relatively good estimate of

³<https://github.com/sz612866/MVKM-Multiview-Tensor>

⁴The code is from <https://github.com/CAHLR/pyBKT>

the real score. As a result, outperforming a simple average baseline is a challenging task.

The mentioned baselines all work on one type of learning material. Since our proposed MVKM model works with more than one learning material type, to be fair in evaluations, we run baseline algorithms in a multi-view setup. To do this, we aggregate the data from all learning material types and use that as an input to these baselines. In those cases, we add a "MV" to the end of their names. For example, FDTF_MV represents running FDTF on aggregation of student interactions with multiple learning material types. In addition, for knowledge tracing algorithms (BKT and DKT) which are designed for binary student responses (correct or incorrect), we modify their settings to make them predict numerical scores as described below. First, we binarize students' historical scores based on median score. Specifically, if the score is greater than the median, it will be set to 1, and 0 otherwise. Then, we use the probability of success generated by BKT and DKT as the probability of student receiving a score more than median score. Eventually, the numerical predicted scores can be obtained by viewing the probability output as the percentile of students' score on that specific question. Moreover, since these models require pre-defined knowledge components (KCs), we assume that each learning material is mapped to one KC in these models.

In addition to the above, we compare our multi-view model with its basic variation (MVKM-Base) using the data from graded materials only, and its multi-view variation without the learning and forgetting constraints (MVKM-W/O-P).

4.3.2 Performance Metrics and Comparison

In this task, our target is to accurately estimate the actual student scores. To evaluate how close our predicted values are to the actual ones, we use Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) between the predicted scores and the actual scores for students. Table 2 and 3 show the results of performance among different methods on synthetic data and real data, respectively. We can see that our proposed model outperforms other baselines on synthetic data, and has the best performance on real datasets in general.

MVKM-Base vs. Single Material Type Baselines.

Comparing MVKM-Base with other algorithms that use student scores only, shows us that MVKM-Base has consistently lower error compared to most baselines, in both synthetic and real-world datasets. This result demonstrates the ability of MVKM-Base in capturing data variance and validity of its assumptions for real-world graded data. Compared to AVG, MVKM-Base can represent more variability; compared to RBTF, the student latent features in MVKM-Base leads to improved results; compared to FDTF, the forgetting factor results in less error; and compared to BKT and DKT, modeling the learning material concepts in Q and having a rank-based constraint to enforce learning improves the performance. The only baseline algorithm that outperforms MVKM-Base in some setups is BPTF. Particularly, BPTF has a lower RMSE and MAE in Synthetic_NG and Synthetic_G datasets that are skewed. In real-world datasets, it performs better than MVKM-Base in MORF-QD dataset that is more sparse and has a slightly higher average score

Methods	Synthetic_NG		Synthetic_NG2		Synthetic_G	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
AVG	0.3084±0.0072	0.2820±0.0093	0.5059±0.0115	0.4005±0.0115	0.3070±0.0039	0.2811±0.0050
RBTF	0.2515±0.0126	0.2027±0.0081	0.3374±0.0234	0.2681±0.0146	0.2628±0.0113	0.2103±0.0080
FDTF	0.4906±0.0172	0.4410±0.0207	0.6588±0.0215	0.5529±0.0226	0.5041±0.0184	0.4537±0.0213
TFWL	0.5283±0.0168	0.4632±0.0178	0.6919±0.0132	0.5883±0.0156	0.5490±0.0053	0.5130±0.0076
BPTF	0.1675±0.0048	0.1256±0.0061	0.3454±0.0140	0.2589±0.0072	0.1825±0.0064	0.1381±0.0050
IBKT	0.4744±0.0118	0.4197±0.0140	0.6630±0.0122	0.5494±0.0152	0.4748±0.0076	0.4233±0.0098
DKT	0.2694±0.0275	0.1911±0.0241	0.4536±0.0404	0.3569±0.0413	0.2716±0.0209	0.2047±0.0178
RBTF-MV	0.2920±0.0069	0.2305±0.0078	0.4064±0.0213	0.3227±0.0147	0.2618±0.0155	0.2126±0.0130
FDTF-MV	0.4078±0.0168	0.3402±0.0167	0.5861±0.0211	0.4688±0.0135	0.4888±0.0112	0.4538±0.0131
TFWL-MV	0.4337±0.0139	0.3896±0.0133	0.6386±0.0161	0.5450±0.0194	0.5312±0.0137	0.4626±0.0145
BPTF-MV	0.1718±0.0037	0.1457±0.0055	0.3438±0.0158	0.2603±0.0120	0.1533±0.0055	0.1184±0.0044
IBKT-MV	0.4257±0.0142	0.3585±0.0155	0.6019±0.0124	0.4892±0.0165	0.4844±0.0068	0.4275±0.0089
DKT-MV	0.4278±0.0313	0.3613±0.0318	0.6399±0.0515	0.5320±0.0526	0.3390±0.0252	0.2892±0.0245
MVKM-Base	0.2007±0.1069	0.1498±0.0809	0.3026±0.0697	0.2273±0.0356	0.2097±0.0485	0.1565±0.0348
MVKM-W/O-P	0.1714±0.0089	0.1306±0.0089	0.2817±0.0316	0.2213±0.0245	0.1796±0.0345	0.1357±0.0190
Our Method (MVKM)	0.1388±0.0048	0.1049±0.0056	0.2221±0.0074	0.1739±0.0048	0.1532±0.0128	0.1171±0.0097

Table 2: Performance Prediction results on synthetic datasets, measured by RMSE and MAE, shown with variance in 5-fold cross-validation

Methods	MORF_QD		MORF_QL		CANVAS_H	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
AVG	0.2410±0.0227	0.1913±0.0161	0.2420±0.0108	0.1957±0.0067	0.0767±0.0121	0.0555±0.0040
RBTF	0.2711±0.0229	0.2132±0.0147	0.2572±0.0114	0.1980±0.0074	0.1571±0.0172	0.1235±0.0103
FDTF	0.3081±0.0437	0.2401±0.0329	0.3006±0.0194	0.2324±0.0151	0.1395±0.0259	0.0929±0.0119
TFWL	0.2750±0.0529	0.2003±0.0249	0.3090±0.3090	0.2237±0.0099	0.2377±0.0803	0.1186±0.0513
BPTF	0.2172±0.0128	0.1776±0.0082	0.2302±0.0068	0.1953±0.0048	0.1114±0.0120	0.0946±0.0082
IBKT	0.2756±0.0070	0.2281±0.0053	0.2646±0.0147	0.2174±0.0096	0.0856±0.0105	0.0692±0.0042
DKT	0.3169±0.0374	0.2498±0.0313	0.2859±0.0061	0.2158±0.0075	0.0911±0.0322	0.0616±0.0173
RBTF-MV	0.2814±0.0282	0.2177±0.0222	0.2624±0.0193	0.1977±0.0136	0.1484±0.0098	0.1171±0.0054
FDTF-MV	0.3138±0.0441	0.2453±0.0387	0.2398±0.0137	0.1866±0.0091	0.1149±0.0085	0.0907±0.0068
TFWL-MV	0.2919±0.0275	0.1975±0.0160	0.3222±0.0208	0.2178±0.0165	0.1748±0.0600	0.0784±0.0269
BPTF-MV	0.2615±0.0129	0.2286±0.0114	0.2313±0.0070	0.1960±0.0041	0.1452±0.0100	0.1343±0.0081
IBKT-MV	0.2774±0.0204	0.2177±0.0099	0.2904±0.0098	0.2137±0.0062	0.0834±0.0125	0.0425±0.0049
DKT-MV	0.2938±0.0310	0.2352±0.0236	0.2540±0.0065	0.2185±0.0047	0.079±0.0247	0.0496±0.0065
MVKM-Base	0.2242±0.0328	0.1669±0.0207	0.2277±0.0119	0.1724±0.0081	0.0666±0.0159	0.0411±0.0040
MVKM-W/O-P	0.2385±0.0196	0.1771±0.0104	0.2450±0.0145	0.1814±0.009	0.0649±0.0111	0.0388±0.0027
Our Method (MVKM)	0.2088 ± 0.0229	0.1603±0.0142	0.2150±0.0127	0.1654±0.0104	0.0613±0.0112	0.0362±0.0028

Table 3: Performance Prediction results on real-world datasets, measured by RMSE and MAE, shown with variance in 5-fold cross-validation.

compared to the other two. This shows that BPTF is better than MVKM-Base in handling skewed data. One potential reason is BPTF’s smoothing assumption, in contrast with MVKM-Base’s rank-based knowledge increase, that results in a more homogeneous score predictions for each student.

MVKM: Multiple vs. Single Material Types. Comparing MVKM’s results with MVKM-Base model, we can see that using data from multiple learning material types improves performance prediction results. It verifies our assumptions regarding knowledge transfer in different learning material types through the knowledge gain in shared concept latent space. This is given that in other models, e.g., all models except DKT in MORF-QD, adding different learning material types increases the prediction error. This error increase is particularly happening with BPTF model in real-world datasets and DKT model in synthetic ones. This shows that merely aggregating data from various resources, without appropriate modeling, can even harm the prediction results. This difference between MVKM and other baselines is in its specific setup, in which each learning material type is modeled separately, while keeping a shared knowledge space, student latent features, and knowledge gain.

Learning and Forgetting Effect. To further test the effect of our knowledge gain and forgetting constraint, we compare MVKM with MVKM-W/O-P, a variation of our

proposed model without the rank-based constraint in Equation 2. We can see that MVKM outperforms MVKM-W/O-P in all datasets. This shows that the soft knowledge increase and forgetting assumption is essential in correctly capturing the variability in students’ learning. Particularly, comparing MVKM-W/O-P’s results with MVKM-Base, the single-view version that includes the rank-based learning constraints, we can measure the effect of adding multiple learning material types vs. the effect of adding the learning and forgetting constraints in MVKM model. In CANVAS.H dataset, adding multiple learning material types is more effective than learning constraint, and in MORF datasets, realizing learning constraint is more important than modeling multiple types of learning materials. Nevertheless, they are not mutually exclusive and both are important in the model.

Hyper-parameter Tuning Using a separate validation set, we experiment with various values (grid search) for model hyper-parameters to select the most representative ones for our data. Specifically, we first vary the student latent feature dimension K in $[1, 5, 10, \dots, 40, 45]$, the question latent feature dimension C in $[1, 2, \dots, 9, 10]$, the penalty weight ω in $[0.01, 0.05, 0.1, 0.5, 1, 2, 3]$, the Markovian step m in $[1, 2, \dots, 10]$, and the learning resource importance parameter $\gamma^{[r]}$ in $[0.05, 0.1, 0.2, 0.5, 1, 2]$. Once we found a good set of hyper-parameters from coarse-grained grid search, we search the values close to the optimal values to find out the

best fine-grained values for these hyper-parameters. The best resulting hyper-parameter values for each dataset are listed in table 4. We use $\gamma^{[1]}$ as the trade-off parameter for graded learning material, $\gamma^{[2]}$ for another learning material. As we can see, in both the synthetic and real-world data, the learning and forgetting constraint is more important (larger ω) when having a non-graded learning material type. This shows that binary interaction data, unlike student grades (or scores), is not precise enough to represent the students' gradual knowledge gain in the absence of a learning and forgetting constraint. Also, comparing $\gamma^{[2]}$ in MORF_QD vs. MORF_QL we can see that the importance of video lectures is more than discussions in predicting students' performance.

Dataset	K	C	ω	$\gamma^{[1]}$	$\gamma^{[2]}$	η	m	λ_t	λ_s
Synthetic_NG	3	3	0.2	1	0.1	0.1	1	0.01	0.001
Synthetic_NG2	3	3	0.2	1	0.1	0.1	1	0.001	0.001
Synthetic_G	3	3	0.1	1	0.4	0.1	1	0.001	0.001
MORF_QD	39	5	1	1	0.05	0.1	1	0	0
MORF_QL	35	9	0.6	1	0.5	0.1	1	0	0
Canvas_H	28	7	2.0	1	0.5	0.01	1	0	0

Table 4: Hyperparameters of our model for each dataset

4.4 Student Knowledge Modeling

In this set of experiments, we answer two main research questions: 1) Can our model's learning and forgetting constraint capture meaningful knowledge trends across concepts for students as a whole? and 2) Are the individual student's knowledge growth representative of their learning? To answer these questions, we look at the estimated knowledge tensor of students ($\mathbf{K} = \mathbf{ST}$).

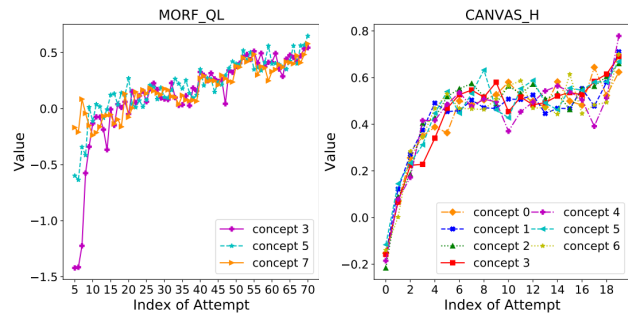


Figure 3: Average knowledge gain of concepts across all students.

To answer the first question, we check the average student knowledge growth on different concepts. Figure 3 shows the average knowledge of all students in different concepts (represented with different colors) during the whole course period (X-axis) for MORF_QL, and CANVAS_H datasets (MORF_QD has similar patterns as MORF_QL, we don't show it due to the page limitation). Notice that, for a clear visualization, we only show 3 out of 9 concepts from MORF_QL dataset in the figure. We can see that, on average, students' knowledge in different concepts increase. Particularly, in MORF_QL, the initial average knowledge on concept 3 is less than concepts 5 and 7. However, students learn this concept rapidly as shown by the increase of knowledge level around the tenth attempt. As the knowledge growth is less smooth in this concept, compared to the other two (e.g., the drop around the 15th attempt), students are more likely to forget it rapidly. Eventually, the average student knowledge in all concepts are close to each other. On

the other hand, in CANVAS_H, the average initial knowledge in different concepts are relatively close. However, students end up having different knowledge levels in different concepts at the end of the course, especially in concepts 0 and 4. Also, all six concepts show large fluctuations across the attempts. Overall, the students have a significant knowledge gain at the first few attempts and the knowledge gain slows down after that. This is aligned with our expectation on students' knowledge acquisition through out the course.

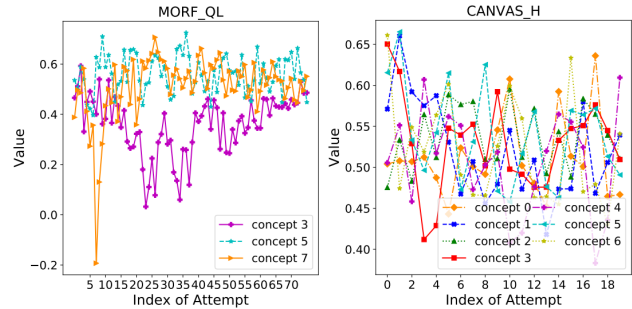


Figure 4: Average knowledge gain of each concept across all students.

To show the effect of the learning and forgetting constraint in MVKM, we look at the student knowledge acquisition in the MVKM-W/O-P model, that removes this constraint. The MVKM-W/O-P's average student knowledge in different concepts throughout all attempts is shown in Figure 4. We can see that despite its acceptable performance prediction error, MVKM-W/O-P's estimated knowledge trends are elusive and counter-intuitive. For example, many concepts (such as concept 3 in MORF_QL) show a U-shaped curve. This curve can be interpreted as the students having a high prior knowledge in these concepts, but forgetting them in the middle of the course, and then re-learning them at the end of the course. In some cases, such as concept 1 in CANVAS_H, students lose some knowledge and forget what they already knew, by the end of the course. This demonstrates the necessity of learning and forgetting penalty term in MVKM.

For second question, we check if there are meaningful differences between knowledge gain trends of different students. To do this, we apply spectral clustering on students' latent features matrix S to discover different groups of students. Then, we compare students' learning curves from different clusters. The number of clusters is determined by the significance of difference on average performance in each cluster. We obtained 3 clusters of students for MORF_QD course, and 2 clusters for MORF_QL and CANVAS_H courses based on students' latent features from our model.

To see the differences in these groups, we sample one student from each cluster in each real-world dataset. Figure 5 shows these sample students' knowledge gain, averaged over all concepts, in datasets MORF_QD and MORF_QL (CANVAS_H is not showed due to the page limitation, it has similar patterns as MORF_QD). The figures show that different students start with different initial prior knowledge. For example, in MORF_QL, student #5 starts with a lower prior knowledge than student #100 and ends up with a lower final knowledge. Also, the figure shows that different knowledge gain trends across students, particularly in MORF_QD. For

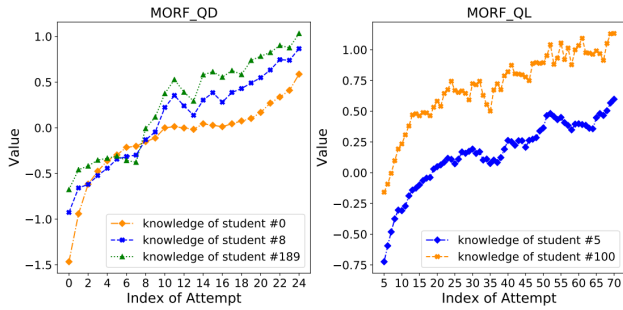


Figure 5: Sample students' knowledge gain across all concepts in two different courses.

example, student #0 starts with a lower prior knowledge than the other two students, but has a faster knowledge growth, and catches up with them around attempt 8. However, this student's knowledge growth slows down after a while and ends up to be lower than the other two at the end of course. To see if the quantified knowledge is meaningful, we compare student's knowledge growth with their scores. Students #0, #8, and #189 in MORF_QD have average grades 0.202, 0.636, and 0.909, in MORF_QL, #5 and #100 have average grades 0.9 and 0.98. This aligns with the knowledge levels shown in the figure. These observations show that MVKM can meaningfully differentiate between different students' knowledge growth.

4.5 Learning Resource Modeling

In this section, we evaluate our model on how well it can represent the variability and similarity of different learning materials. We mainly focus on two questions: 1) Are the learning materials' biases consistent with their difficulty levels? 2) Are the discovered latent concepts for learning materials (matrix $Q^{[r]}$) representative of actual conceptual groupings of learning materials in the real datasets?

Bias Evaluation. For the first question, since we do not have access to the learning materials' difficulty levels, we use average student scores on them, as a proxy for difficulty. As a result, we only use graded learning materials for this analysis. We calculated the spearman correlation between question bias captured by our model and average score of each question. The spearman correlation on MORF_QD is 0.779, on MORF_QL is 0.597, and on CANVAS_H is 0.960. We find that question bias derived from MCKM is highly correlated with average question score, where the lower the actual average grades are, the lower the bias values are learned.

Within-Type Concept Evaluation. For the second question, we would like to know how much the learning materials' discovered latent concepts resemble the real-world similarities in them. To evaluate the real-world similarities, we rely on two scenarios: 1) the learning material that are arranged closely to each other in the course structure, either in same module or in consequent modules, are similar to each other (course structure similarity); 2) the learning materials that are similar to each other have similar concepts and contents (content similarity). Since only one of our real-world datasets, MORF_QL, includes the required information for these scenarios, we use this dataset in the continuation of this paper. For first scenario, the course includes an ordered list modules, each of which include an ordered list of videos,

in addition to the assignments associated with each module.

For the second scenario, because our learning materials are not labeled with their concepts in our datasets, we use their textual contents (not used in MVKM) as a representation of their concepts. Particularly, we have subtitles for 40 video lectures, and text of questions for 8 quizzes. We note that if two learning materials present the same concepts, their textual contents should also be similar to each other. As a result, we build content-based clusters of learning materials, each of which containing the learning materials that are conceptually similar to each other. Specifically, to cluster the learning material according to their contents, we use Spectral Clustering on the latent topics that are discovered using Latent Dirichlet Analysis (LDA)[9] on the learning material's textual contents. In the same way, we can cluster the learning materials according to their discovered latent concepts by MVKM. Similar to the textual analysis, we use spectral clustering on the discovered $Q^{[r]}$ matrices to form clusters of learning materials. To do this, we first consider only one learning material type (the video lectures) and then move on to the similarities between two types of learning materials (both video lectures and assignments).

The results are shown in Figure 6 for within-type learning material similarity in video-lectures. Figure 6(a) shows the 8 clusters that were discovered using MVKM, and Figure 6(b) shows the 8 clusters that were discovered using video-lecture transcripts. Each cluster is shown within a box with a number associated with it. Each video-lecture is shown by its module (or week in the course), its order in the module sequence, and its name. For ease of comparison, we colored the video names according to their LDA content clusters. Looking at the LDA content clusters, we can see that although some lectures in same module fit in same cluster (e.g., videos 1, 2, 3, and 4 from week 7 are all in cluster 7), some of the lectures do not cluster with other videos in their module. For example, video 5 in week 7 is in cluster 2, with pioneer knowledge tracing methods. This shows that in addition to structural similarities, content similarities also exist in learning materials. Looking at MVKM clusters, we can see that the clusters mostly represent the course structure similarity: learning materials from same module are grouped. For example, all videos of week 3 are grouped in cluster 2. However, we can see that in many cases, whenever the structure similarity in clusters are disrupted, it is because of the content similarity in video lectures. For example, video 5 in week 7 that was clustered with pioneer knowledge tracing method in LDA content clusters is also clustered with them in MVKM clusters.

Between-Type Concept Evaluation. To evaluate MVKM's discovered similarities between different types of learning materials, we evaluate assignments' and video lectures' in MORF_QL. To do this, we build LDA-based clusters using assignment texts and video lecture transcripts. These clusters are shown in Figure 7(b). We also cluster the learning materials using spectral clustering on the concatenation of their $Q^{[r]}$ matrices (Figure 7(a)). Because the assignments bring more information to the clustering algorithms, the clustering results are different from the clusters of video lectures only. Similar to within-type concept evaluation results, we can still see the effect of both content and structure

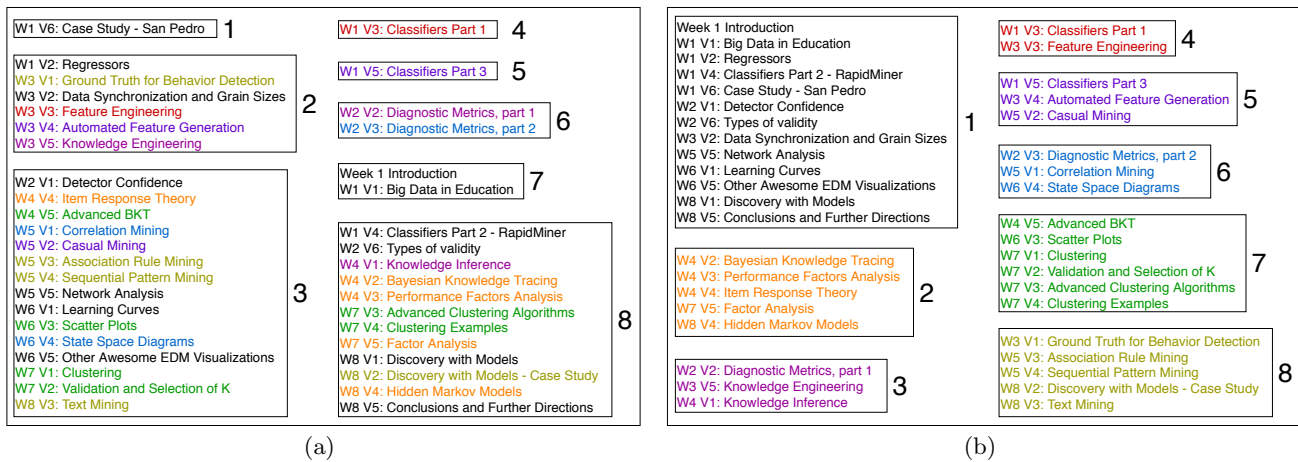


Figure 6: Clusters that were discovered by using MVKM (a), clusters discovered by using video-lecture transcripts (b).

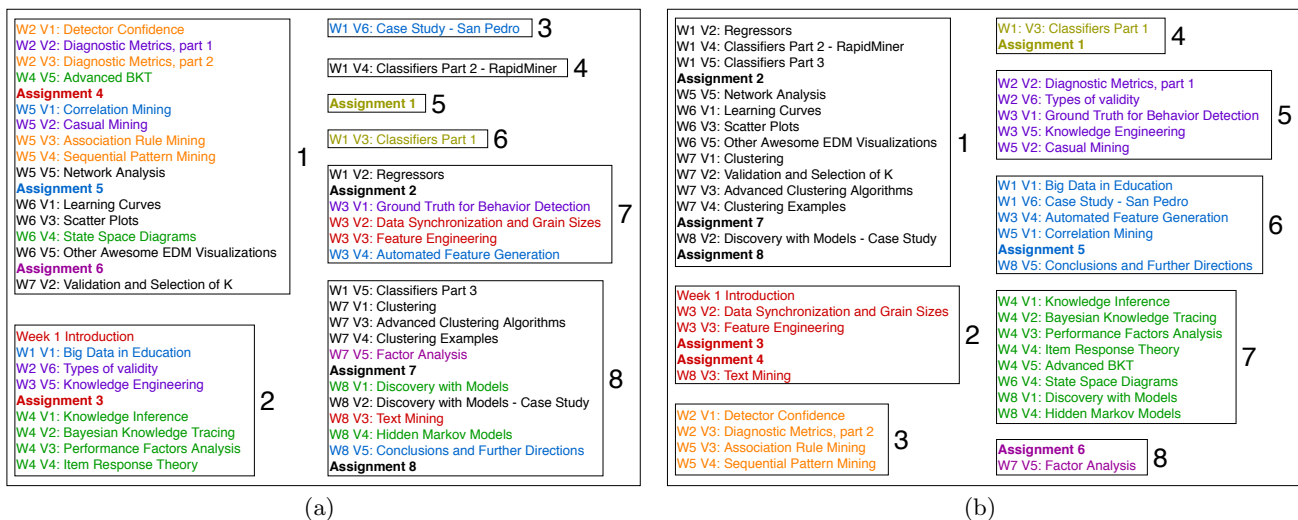


Figure 7: Clusters discovered by using MVKM (a), clusters discovered by using video-lecture transcripts and assignment texts(b).

similarities in video lectures that are clustered together by MVKM. For example, videos 1 and 3 of week 2 are clustered with later weeks’ videos because of content similarity (cluster 1 in Figure 7(a)). While videos 2 of week 2 is also clustered with them because it comes between these two videos in course sequence.

Additionally, between video lectures and assignments, the clusters closely follow the course structure. The assignments in this course come at the end of their module and right before the next module starts. For example, “Assignment 3” appears after video 5 at week 3 and before video 1 at week 4. We can see that all assignments, except “Assignment 1” that is the first one, are clustered with their immediate next video lecture. Moreover, we can see the effect of content similarity between assignments and video lectures in differences of Figures 6(a) and 7(a). For example, without including assignments, “Week 1 Introduction” and “W1 V1: Big Data in Education” were clustered together in cluster 7 of Figure 6(a). However, after adding assignments, because of the content similarity between “Assignment 3” and “Week 1 Introduction” (Figure 7(b) cluster 2), “Week 1 Introduction” and “W1 V1: Big Data in Education” are clustered with video lectures that are structurally close to “Assignment 3”.

Altogether, we demonstrated that learning materials’ bias parameters in MVKM are aligned with their difficulties; learning materials’ latent concepts discovered by our model well represent learning materials’ real-world similarities, both in structure and in content; and MVKM can successfully unveil these similarities between different types of learning materials, without observing their content or structure.

5. CONCLUSIONS

In this paper, we proposed a novel Multi-View Knowledge Model (MVKM) that can model students’ knowledge gain from different learning materials types, while simultaneously discovering materials’ latent concepts. Our proposed tensor factorization model explicitly represents students’ knowledge growth and allows for occasional forgetting of learned concepts. Our extensive evaluations on synthetic and real-world datasets show that MVKM outperforms other baselines in the task of student performance prediction, can effectively capture students’ knowledge growth, and represent similarities between different learning materials types.

6. ACKNOWLEDGMENTS

This paper is based upon work supported by the National Science Foundation under Grant No. 1755910.

7. REFERENCES

- [1] R. Agrawal, M. Christoforaki, S. Gollapudi, A. Kannan, K. Kenthapadi, and A. Swaminathan. Mining videos from the web for electronic textbooks. In *International Conference on Formal Concept Analysis*, pages 219–234. Springer, 2014.
- [2] R. Agrawal, S. Gollapudi, A. Kannan, and K. Kenthapadi. Enriching textbooks with images. In *Proceedings of the 20th ACM International Conference on Information and knowledge management*, pages 1847–1856, 2011.
- [3] G. Alexandron, Q. Zhou, and D. Pritchard. Discovering the pedagogical resources that assist students in answering questions correctly—a machine learning approach. In *the 8th International Conference on Educational Data Mining (EDM 2015)*, pages 520–523, 2015.
- [4] J. M. L. Andres, R. S. Baker, G. Siemens, D. Gašević, and C. A. Spann. Replicating 21 findings on student success in online learning. *Technology, Instruction, Cognition, and Learning*, pages 313–333, 2016.
- [5] R. S. Baker. *Big Data and Education*. New York, NY: Teachers College, Columbia University, 2nd edition, 2015.
- [6] R. S. Baker, A. T. Corbett, and V. Aleven. More accurate student modeling through contextual estimation of slip and guess probabilities in bayesian knowledge tracing. In *International conference on intelligent tutoring systems*, pages 406–415. Springer, 2008.
- [7] T. Barnes. The q-matrix method: Mining student response data for knowledge. In *American Association for Artificial Intelligence 2005 Educational Data Mining Workshop*, pages 1–8, 2005.
- [8] J. E. Beck, K.-m. Chang, J. Mostow, and A. Corbett. Does help help? introducing the bayesian evaluation and assessment methodology. In *International Conference on Intelligent Tutoring Systems*, pages 383–394. Springer, 2008.
- [9] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- [10] A. F. Botelho, S. Adjei, and N. T. Heffernan. Modeling interactions across skills: A method to construct and compare models predicting the existence of skill relationships. In *EDM*, pages 292–297, 2016.
- [11] Canvas-Network. Canvas network courses, activities, and users (4/2014 - 9/2015) restricted dataset, 2016.
- [12] G. Casalino, C. Castiello, N. Del Buono, F. Esposito, and C. Mencar. Q-matrix extraction from real response data using nonnegative matrix factorizations. In *International Conference on Computational Science and Its Applications*, pages 203–216. Springer, 2017.
- [13] Y. Chen, J. P. González-Brenes, and J. Tian. Joint discovery of skill prerequisite graphs and student models. In *The 9th International Educational Data Mining*, pages 46–53, 2016.
- [14] B. Choffin, F. Popineau, Y. Bourda, and J.-J. Vie. Das3h: Modeling student learning and forgetting for optimally scheduling distributed practice of skills. In *The 12th International Conference on Educational Data Mining*, pages 384–389. IEDMS, 2019.
- [15] A. T. Corbett and J. R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4(4):253–278, 1994.
- [16] M. Desmarais, P. Xu, and B. Beheshti. Combining techniques to refine item to skills q-matrices with a partition tree. In O. Santos, J. G. Boticario, C. Romero, M. Pechenizkiy, A. Merceron, P. Mitros, J. M. Luna, C. Mihaescu, P. Moreno, A. Hershkovitz, S. Ventura, and M. Desmarais, editors, *the 8th International Conference on Educational Data Mining (EDM 2015)*, pages 29–36, 2015.
- [17] T.-N. Doan and S. Sahebi. Rank-based tensor factorization for student performance prediction. In *Proceedings of The 12th International Conference on Educational Data Mining (EDM 2019)*, volume 288, page 293. ERIC.
- [18] T.-N. Doan and S. Sahebi. Rank-based tensor factorization for predicting student performance. In *The 12th International Conference on Educational Data Mining*, pages 288–293. IEDMS, 2019.
- [19] F. Drasgow and C. L. Hulin. Item response theory. *Handbook of industrial and organizational psychology*, 1:577–636, 1990.
- [20] J. González-Brenes. Modeling skill acquisition over time with sequence and topic modeling. In *Artificial Intelligence and Statistics*, pages 296–305, 2015.
- [21] J. P. González-Brenes, Y. Huang, and P. Brusilovsky. Fast: Feature-aware student knowledge tracing. In *NIPS Workshop on Data Driven Education*, 2013.
- [22] R. Hosseini, T. Sirkiä, J. Guerra, P. Brusilovsky, and L. Malmi. Animated examples as practice content in a java programming course. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education, SIGCSE '16*, pages 540–545, New York, NY, USA, 2016. ACM.
- [23] Y. Huang, J. P. González-Brenes, and P. Brusilovsky. Challenges of using observational data to determine the importance of example usage. In *International Conference on Artificial Intelligence in Education*, pages 633–637. Springer, 2015.
- [24] M. J. Johnson. Scaling cognitive modeling to massive open environments.
- [25] M. M. Khajah, Y. Huang, J. P. González-Brenes, M. C. Mozer, and P. Brusilovsky. Integrating knowledge tracing and item response theory: A tale of two frameworks. In *Proceedings of Workshop on Personalization Approaches in Learning Environments (PALE 2014) at the 22th International Conference on User Modeling, Adaptation, and Personalization*, pages 7–12. University of Pittsburgh, 2014.
- [26] K. R. Koedinger, J. C. Stamper, E. A. McLaughlin, and T. Nixon. Using data-driven discovery of better student models to improve student learning. In *Artificial intelligence in education*, pages 421–430. Springer, 2013.
- [27] A. S. Lan, C. Studer, and R. G. Baraniuk. Time-varying learning and content analytics via sparse factor analysis. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 452–461, 2014.
- [28] A. S. Lan, C. Studer, and R. G. Baraniuk.

- Time-varying learning and content analytics via sparse factor analysis. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 452–461. ACM, 2014.
- [29] R. V. Lindsey, J. D. Shroyer, H. Pashler, and M. C. Mozer. Improving students’ long-term knowledge retention through personalized review. *Psychological science*, 25(3):639–647, 2014.
- [30] R. Liu, J. L. Davenport, and J. C. Stamper. Beyond log files: Using multi-modal data streams towards data-driven kc model improvement. In *The 9th International Educational Data Mining*, pages 436–441, 2016.
- [31] M. C. Mozer and R. V. Lindsey. Predicting and improving memory retention: Psychological theory matters in the big data era. In *Big data in cognitive science*, pages 43–73. Psychology Press, 2016.
- [32] K. Nagatani, Q. Zhang, M. Sato, Y.-Y. Chen, F. Chen, and T. Ohkuma. Augmenting knowledge tracing by considering forgetting behavior. In *The World Wide Web Conference*, pages 3101–3107. ACM, 2019.
- [33] A. S. Najjar, A. Mitrovic, and B. M. McLaren. Adaptive support versus alternating worked examples and tutored problems: Which leads to better learning? In *International Conference on User Modeling, Adaptation, and Personalization*, pages 171–182. Springer, 2014.
- [34] Z. Pardos, Y. Bergner, D. Seaton, and D. Pritchard. Adapting bayesian knowledge tracing to a massive open online course in edx. In *Educational Data Mining 2013*, 2013.
- [35] Z. A. Pardos and N. T. Heffernan. Modeling individualization in a bayesian networks implementation of knowledge tracing. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 255–266. Springer, 2010.
- [36] P. I. Pavlik, H. Cen, and K. R. Koedinger. Performance factors analysis-a new alternative to knowledge tracing. In *14th International Conference on Artificial Intelligence in Education*, volume 2009, pages 531–538, 2009.
- [37] R. Pelánek. Measuring similarity of educational items: An overview. *To appear in IEEE Transactions on Learning Technologies*, 2019.
- [38] C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein. Deep knowledge tracing. In *Advances in neural information processing systems*, pages 505–513, 2015.
- [39] Y. Qiu, Y. Qi, H. Lu, Z. A. Pardos, and N. T. Heffernan. Does time matter? modeling the effect of time with bayesian knowledge tracing. In *EDM*, pages 139–148, 2011.
- [40] S. Sahebi. *Canonical Correlation Analysis in Cross-Domain Recommender Systems*. PhD thesis, Intelligent Systems Program, University of Pittsburgh, 2016.
- [41] S. Sahebi and P. Brusilovsky. Student performance prediction by discovering inter-activity relations. *Educational Data Mining*, pages 87–96, 2018.
- [42] S. Sahebi, Y. Lin, and P. Brusilovsky. Tensor factorization for student modeling and performance prediction in unstructured domain. In *The 9th International Conference on Educational Data Mining*, pages 502–506. IEDMS, 2016.
- [43] M. Sao Pedro, R. Baker, and J. Gobert. Incorporating scaffolding and tutor context into bayesian knowledge tracing to predict inquiry skill acquisition. In *Educational Data Mining 2013*, 2013.
- [44] N. Thai-Nghe, T. Horvath, and L. Schmidt-Thieme. Context-aware factorization for personalized student’s task recommendation. In *Proceedings of the International Workshop on Personalization Approaches in Learning Environments*, volume 732, pages 13–18, 2011.
- [45] W. J. van der Linden and R. K. Hambleton. *Handbook of modern item response theory*. Springer Science & Business Media, 2013.
- [46] N. F. Velasquez, I. Goldin, T. Martin, and J. Maughan. Learning aid use patterns and their impact on exam performance in online developmental mathematics. In *Educational Data Mining 2014*, 2014.
- [47] J.-J. Vie and H. Kashima. Knowledge tracing machines: Factorization machines for knowledge tracing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 750–757, 2019.
- [48] L. Wang, A. Sy, L. Liu, and C. Piech. Learning to represent student knowledge on programming exercises using deep learning. In *Proceedings of the 10th International Conference on Educational Data Mining*, pages 324–329, 2017.
- [49] M. Wen and C. P. Rosé. Identifying latent study habits by mining learner behavior patterns in massive open online courses. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 1983–1986. ACM, 2014.
- [50] L. Xiong, X. Chen, T.-K. Huang, J. G. Schneider, and J. G. Carbonell. Temporal collaborative filtering with bayesian probabilistic tensor factorization. In *SIAM International Conference on Data Mining*, volume 10, pages 211–222, 2010.
- [51] M. V. Yudelson, K. R. Koedinger, and G. J. Gordon. Individualized bayesian knowledge tracing models. In *International Conference on Artificial Intelligence in Education*, pages 171–180. Springer, 2013.
- [52] M. V. Yudelson, O. P. Medvedeva, and R. S. Crowley. A multifactor approach to student model evaluation. *User Modeling and User-Adapted Interaction*, 18(4):349–382, 2008.
- [53] J. Zhang, X. Shi, I. King, and D.-Y. Yeung. Dynamic key-value memory networks for knowledge tracing. In *Proceedings of the 26th international conference on World Wide Web*, pages 765–774. International World Wide Web Conferences Steering Committee, 2017.