Learning Skill Equivalencies Across Platform Taxonomies

Zhi Li

University of California, Berkeley
zhili@berkeley.edu

Xianyou Li

University of California, Berkeley
lixianyou@berkeley.edu

Cheng Ren

University of California, Berkeley
cheng.ren@berkeley.edu

Zachary A. Pardos

University of California, Berkeley
pardos@berkeley.edu

ABSTRACT

Assessment and reporting of skills is a central feature of many digital learning platforms. With students often using multiple platforms, cross-platform assessment has emerged as a new challenge. While technologies such as Learning Tools Interoperability (LTI) have enabled communication between platforms, reconciling the different skill taxonomies they employ has not been solved at scale. In this paper, we introduce and evaluate a methodology for finding and linking equivalent skills between platforms by utilizing problem content as well as the platform's clickstream data. We propose six models to represent skills as continuous real-valued vectors, and leverage machine translation to map between skill spaces. The methods are tested on three digital learning platforms: ASSISTments, Khan Academy, and Cognitive Tutor. Our results demonstrate reasonable accuracy in skill equivalency prediction from a fine-grained taxonomy to a coarse-grained one, achieving an average recall@5 of 0.8 between the three platforms. Our skill translation approach has implications for aiding in the tedious, manual process of taxonomy to taxonomy mapping work, also called crosswalks, within the tutoring as well as standardized testing worlds.

CCS CONCEPTS

• **Applied computing** → **Education**; • **Computing methodologies** → **Learning latent representations**; **Machine translation**.

KEYWORDS

Skill equivalencies, transfer models, crosswalks, taxonomies, digital learning platforms, representation learning, machine translation, interoperability, acknowledging prior knowledge, app hand-offs.

ACM Reference Format:

Zhi Li, Cheng Ren, Xianyou Li, and Zachary A. Pardos. 2021. Learning Skill Equivalencies Across Platform Taxonomies. In *LAK21: 11th International Learning Analytics and Knowledge Conference (LAK21)*, April 12–16, 2021, Irvine, CA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3448139.3448173>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LAK21, April 12–16, 2021, Irvine, CA, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8935-8/21/04...\$15.00

<https://doi.org/10.1145/3448139.3448173>

1 INTRODUCTION

Digital learning platforms commonly tag skills to assessment items in order to measure students' mastery and guide their learning trajectories [3, 14, 29]. Due to the formative nature of digital learning platforms [12], their skill taxonomies are generally finer-grained as compared to those found in large scale summative tests, where broader constructs and abilities are measured [33]. The presence of an accurate skill or knowledge component model [14] in a digital learning platform can have a significant positive impact on its efficacy [4]. Conversely, the absence of an accurate expert skill model can impede effective assessment of mastery and subsequently prohibit adaptive learning approaches [22].

Customarily, digital learning platforms have developed their own taxonomies, but the demand for greater interoperability across platforms is rapidly growing [28]. Teachers and students today often use a mix of devices and many learning platforms per device in class [9]. To effectively assess students and acknowledge prior learning when they switch platforms, smooth "hand-offs" [10] between apps are needed, including a shared taxonomy or translation between taxonomies enabling learning progress on one platform to be continued on the next. Many platforms in the United States have already embarked on migrating to a common taxonomy such as the Common Core State Standards [1] and re-tagging their content or re-mapping their in-house taxonomy to this standard. This task is highly time consuming and likely to be repeated every time a new set of common standards are introduced or modified.

In general, the task of validating the equivalence of skills across taxonomies is non-trivial. For example, the skill of "area of irregular figure" in system A and "area of quadrilaterals and polygons" in system B may or may not be strongly related and it is almost impossible to determine without looking at the set of problems associated with each.

To date, mapping taxonomies across digital learning platforms using machine learning has largely remained unexplored. They are, however, a prime candidate for this type of approach as they offer data affordances distinct from the standardized testing context, where most taxonomy mapping has occurred. In particular, platforms store data about students' learning behaviors such as logged clickstream and response sequence data [11]. Some research has successfully harnessed response data to infer the underlying skill of problems [24]. Here, we evaluate the utility of response sequence data, in addition to problem text, for cross-platform taxonomy mapping.

We show that a taxonomy mapping, or transfer model, can be learned between the skill vector spaces of different platforms, similar to how machine language translation learns mappings between word embedding spaces [17]. Our main contributions are:

- (1) Six models proposed to represent skills as continuous real-valued vectors, with each model able to represent different input data.
- (2) Empirical validation of the feasibility of employing machine translation to map between the taxonomies of three digital learning platforms.
- (3) Further extension of the methodology to situations where two platforms have asymmetric data types available for representation (e.g., one platform with problem text, the other with only response sequences).
- (4) Inspection of important factors that have an influence on skill equivalency prediction performance, including model hyperparameters and differences between source and destination taxonomy granularity.

2 RELATED WORK

Taxonomy mapping has relied on manual work by subject matter experts [7, 29, 32], though there has been past research on automating the process using legacy Natural Language Processing (NLP) to find similar skills across taxonomies using text descriptions of each skill [6, 35]. Choi et al. [6] converted each skill statement to a verb phrase graph and a noun phrase graph, then calculated similarity between skills by comparing graphs. Yilmazel et al. [35] used rule-based methods to extract features from skill descriptions in one standard and trained a machine learning classifier to map them to another standard.

Several terms have been used in the literature to refer to the mappings between skill taxonomies. The term "crosswalk" is one, derived from the idea of creating a path to cross a street, used to describe the connection between two taxonomies or sets of educational standards [7, 32]. Other terms like "transfer" [29] and "alignment" [6, 35] have also been used in related work. In the context of digital learning environments, several terms have been used to refer to the elements of their taxonomy. Intelligent Tutoring Systems refer to each element as a knowledge component (KC) and a taxonomy of elements as a knowledge component model. A KC is defined as, "an acquired unit of cognitive function or structure that can be inferred from performance on a set of related tasks" [14]. These KCs, tagged to problem steps, allow for adaptive tutoring and cognitive mastery estimation in systems like the Cognitive Tutor (now MATHia) [4]. The generic term of "skill" [29] or "tag" [28] has been more commonly used in digital learning environments to describe a semi-granular subject area associated with a learning resource. The term "skill" can also be used to generalize the concepts of tags and knowledge components, which is how we will use it throughout the paper. Differences in granularity and epistemology of taxonomies add to the challenge of cross-platform skill equivalency learning.

Methodologically, no prior work has utilized problem text or response sequences, nor have modern neural approaches from computational linguistics been brought to bear on taxonomy mapping. In the broader field of learning analytics research, neural word

embeddings [18] have been utilized in a number of areas where text descriptions of educational resources are available. Examples include detecting student misconceptions [16], extracting course concepts [20], and coding non-cognitive traits related to student success [31]. Non-text, sequence data can also carry useful semantic information, such as sequences in which an educational resource or skill appears (i.e., its clickstream logs), which we call "context" information. Previous research has shown that item embeddings learned from sequence contexts with a skip-gram model [18] can encode underlying attributes of items helpful to downstream institutional prediction tasks [13], skill label inference [24], and course recommendation [19, 25]. After item embeddings are learned, a translation model can be trained to map items from one embedding to another. This is called machine translation and was introduced in the context of language translation using word embeddings [2, 17]; however, the idea can be extended to translation between any embedding spaces, such as translating course embeddings between institutions to identify candidate courses for credit articulation [23].

3 MODELS AND METHODOLOGY

Each platform consists of problems students are expected to solve that are associated with one or more skills labeled by domain experts. We define the problem of skill equivalency prediction across platforms as follows: given a skill s in a source platform src , find k most similar skills ordered by similarity in destination dst . The input includes the content and/or context information of skills and some ground truth equivalent skill pairs. More precisely, our method is as follows:

- (1) Represent each skill as a continuous real-valued vector using problem content and/or sequence context.
- (2) Learn a translation model from the source skill space to the destination skill space.
- (3) Calculate cosine similarities between skills across platforms:

$$\begin{aligned} \text{cosine_similarity}(s_{src}, s_{dst}) &= \frac{s_{src} \cdot s_{dst}}{\|s_{src}\| \|s_{dst}\|} \\ &= \frac{\sum_{i=1}^n s_{src}^i s_{dst}^i}{\sqrt{\sum_{i=1}^n (s_{src}^i)^2} \sqrt{\sum_{i=1}^n (s_{dst}^i)^2}} \end{aligned}$$

- (4) For each source skill, rank the destination skills by similarity and take the top k as predictions.

In this section, we describe the models for skill representation and translation. Representation models can be categorized into three types: content-based models (Section 3.1), context-based model (Section 3.2), and models combining content and context (Section 3.3). The way in which we will translate one representation to another is borrowed from machine translation (Section 3.4). Code to replicate our methodology can be found online¹.

3.1 Content-based Models

We represent skills as functions of the problem content they are associated with. The content of a problem in a digital learning platform can include text, graphical figures, and video. For this study, our data only contain the text portion of the content. We

¹<https://github.com/CAHLR/skill-equivalency>

utilize the following three content-based representations: Bag-of-words, TF-IDF, and Content2vec.

3.1.1 Bag-of-words. Bag-of-words is a standard text processing technique where a document is represented as a vector whose length is equal to the vocabulary size and values are the frequencies of the words occurring in the document. In our experiments, we take all unique words from both platforms as the vocabulary, and represent each problem as a Bag-of-words vector. Then for every skill, we find all problems associated with the skill and average their representations together arriving at a skill vector.

3.1.2 TF-IDF. A weakness of Bag-of-words is that the vector might be dominated by frequent but non-distinguishing words. TF-IDF (Term Frequency-Inverse Document Frequency), a method adapted from Bag-of-words, can address this issue, where the values now reflect how distinct a word is to a problem relative to the collection of all problems. The TF-IDF score for word w in problem p from problem set P is calculated by:

$$TF\text{-}IDF(w, p, P) = TF(w, p) \cdot IDF(w, P) \quad (1)$$

where

$$TF(w, p) = \log(1 + freq(w, p)) \quad (2)$$

$$IDF(w, P) = \log\left(\frac{|P|}{count(x \in P : w \in x)}\right) \quad (3)$$

3.1.3 Content2vec. Content2vec is a method built on word embeddings. While previous works have mainly utilized word vectors pretrained on large-scale datasets like Google News corpus [16, 20, 31], our problem space is different from theirs in that mathematical jargon, symbols, and formulas that are not frequently seen in other corpora take up the majority of our problem texts. Therefore, we test two versions of Content2vec, one with word vectors pretrained on Google News and the other with our own word vectors trained on all the problem texts. We represent each problem as the average of the word vectors, and each skill as the average of the problem vectors.

Note that in content-based methods, representations of problems and skills across platforms share the same space and thus there is no need to translate between spaces for cross-platform skill comparison.

3.2 Context-based Model (Skill2vec)

Context information comes from the response log data of a digital learning platform, where each row denotes an interaction between a student and the platform, and pertinent columns include anonymized student ID, attempted problem ID, tagged skills, and start time. To make use of context information we propose Skill2vec. In this method, data are preprocessed into skill sequences per student ordered by start time, and a skip-gram model is then applied to these sequences to learn a continuous vector embedding for each skill. The model is similar to previous work in which problems were embedded based on problem sequence [24], but has not until now been applied to skills.

Unlike content-based methods, Skill2vec is trained separately on each platform, hence the generated skill vector spaces are not aligned. Therefore, before measuring cross-platform skill similarity,

a translation is necessary to align the two vector spaces. Details are discussed in Section 3.4.

3.3 Models Combining Content and Context

Since both content and context data may contain useful information about skills, we also evaluate the combination of them that may form a more descriptive representation than either alone. In this section, we introduce two models that combine content and context information: a simple concatenation of our previously mentioned content and context vectors and a separate model from the literature, called Text-Associated Matrix Factorization (TAMF), that integrates both during training.

3.3.1 Content2vec+Skill2vec. The first model is simply the concatenation of Content2vec and Skill2vec. Since Content2vec does not need translation but Skill2vec does, we design the process of concatenation as follows (Figure 1): we first learn a translation model for context vectors, and then combine the translated source context vector with the source content vector, while the destination skill vector is obtained by appending the destination content vector to the unchanged destination context vector. By doing so, we make sure the two components reside in the same space, thus allowing the concatenated vectors to be compared.

Figure 1: Process of combining Skill2vec and Content2vec representations

3.3.2 Text-Associated Matrix Factorization. While concatenation of representations is commonplace, we also sought out a method that learns a single representation from both sources. Text-Associated Matrix Factorization (TAMF) [34] is such a method and may be able to learn features related to the interaction between content and context, and could in principle be more expressive than simple concatenation where no interaction between sources can be utilized.

The method learns embeddings based on nodes in a graph and incorporates content information in a matrix factorization stage. Our adaptation differs from the original model in that they begin with a graph and use Deepwalk [26] to generate a context matrix from which item embeddings are learned, while our response log data are already sequences, therefore we factorize the Positive Pointwise Mutual Information (PPMI) matrix derived from our skill sequences instead of a Deepwalk matrix.

Specifically, let S be the set of all skills and P be the set of all skill and context skill pairs observed in the input data, then the PPMI matrix $M \in \mathbb{R}^{|S| \times |S|}$ can be calculated by $M_{s,c} = \max(0, \log \frac{\#(s,c) \cdot |P|}{\#(s) \cdot \#(c)})$ where s is a skill, c is a context skill, and $\#(s)$, $\#(c)$, and $\#(s,c)$ are the counts of occurrences of s , c , and the pair s, c within a window size in all sequences. According to Levy and Goldberg [15], a skip-gram model (Skill2vec in our case) implicitly factorizes the PPMI matrix $M = W^T H$ where $W, H \in \mathbb{R}^{k \times |S|}$, with k being the dimension of the learned embeddings. The matrix W is equivalent to the output embeddings of the skip-gram model.

With TAMF, the problem is formulated as $M = W^T H T$, where $W \in \mathbb{R}^{k \times |S|}$, $H \in \mathbb{R}^{k \times d}$, and $T \in \mathbb{R}^{d \times |S|}$ (Figure 2). The new matrix T is the Content2vec matrix of embedding size d . The output of the method is the concatenation $[W^T, (HT)^T] \in \mathbb{R}^{|S| \times 2k}$ with embedding size $2k$. In this way content and context are deeply merged in the skill representations.

Figure 2: Text-Associated Matrix Factorization (TAMF)

The optimization objective of TAMF is to search for W and H that minimizes the loss:

$$L = \|M - W^T H T\|_F^2 + \frac{\lambda}{2} (\|W\|_F^2 + \|H\|_F^2)$$

where the first term $\|M - W^T H T\|_F^2$ is the distance between the original PPMI matrix and the reconstructed matrix from the learned factorization, and the second term $\frac{\lambda}{2} (\|W\|_F^2 + \|H\|_F^2)$ is a regularization term restricting the magnitude of the learned matrix. Minimizing this loss should give us a reasonable factorization while not overfitting to the data.

This loss L is convex with respect to W when H is fixed and vice versa, hence we can iteratively optimize in closed form by taking partial derivatives, reorganizing, and solving linear systems for W and H given the other one fixed until the decreasing rate of L is below a certain threshold. While this process does not guarantee to hit the global minimum, in practice we find it works well.

3.4 Machine Translation

For Skill2vec and TAMF, representations of skills in different platforms are learned independently and do not share the same space. They may not even have the same dimensionality. To align vector spaces, we use machine translation to learn a transformation from

a source skill vector space to destination space. Previous research has found that linear transformation outperforms more complex neural translation models on word-to-word translation tasks [17]. Therefore we choose to learn a linear translation matrix $T \in \mathbb{R}^{m \times n}$ that maps a source skill vector $v_s \in \mathbb{R}^n$ to a vector $v_d \in \mathbb{R}^m$ in the destination vector space by maximizing the cosine similarities of translated source vectors and ground truth destination vectors based on a set of known equivalent skill pairs.

4 EXPERIMENTAL SETUP

4.1 Datasets

The datasets for this study are from three digital learning platforms: ASSISTments², Khan Academy³, and Cognitive Tutor (now named MATHia⁴). All three platforms offer content primarily for middle school and high school students.

ASSISTments. We use the public ASSISTments 2012-2013 dataset⁵ including problem texts, which contains problems for 7th and 8th grades math [12]. Most problems in ASSISTments are tagged with a single skill and we keep only these problems and keep only skills with at least 1,000 response logs associated with them. We apply this filter since ASSISTments contains much content that is teacher produced for a single class. ASSISTments has the most coarse-grained taxonomy among the three platforms.

Khan Academy. We use anonymized student responses collected from Khan Academy 7th and 8th grades math exercises between 2013-2014, and collect problem texts through web scraping. Because there are no explicit skills assigned in our data, and each exercise in Khan Academy serves as a template generating multiple problems, we decide to regard each exercise as a skill, as was done in Piech et al. [27]. This results in a very minor difference between some skills, (e.g., *combining_like_terms_1* and *combining_like_terms_2* are considered as two skills). We find that keeping these skills separate gives better performance than grouping them together. Khan Academy's taxonomy is finer than ASSISTments, but coarser than Cognitive Tutor.

Cognitive Tutor. We use the publicly available Cognitive Tutor dataset from the 2010 KDD Cup [30]. We choose the Algebra I 2008-2009 challenge dataset as it is the largest and its content covers 8th-10th grade math [21] with good overlap with the other two platforms in our study. Cognitive Tutor skills are assigned per step, rather than per problem, and it is not uncommon for each step to have more than one skill. To allow for multiple skills, we generate the skill sequences for Skill2vec by randomly ordering skills within a single step. We select the column "KC(SubSkills)" to serve as the skills column since it has the least number of missing values compared to alternative KC columns. Note that problem texts (i.e. content information) are not present in Cognitive Tutor, raising the issue of asymmetric data addressed in Section 5.2. The taxonomy of Cognitive Tutor is the finest-grained.

²<https://www.assistments.org>

³<https://www.khanacademy.org>

⁴<https://www.carnegielearning.com>

⁵<https://sites.google.com/site/assistmentsdata/home/2012-13-school-data-with-affect>

Table 1: Descriptive statistics of datasets (after preprocessing)

	ASSISTments	Khan Academy	Cognitive Tutor
Subject	Math	Math	Math
Grade Level	7th and 8th grades	7th and 8th grades	8th to 10th grades
Number of Skills	130	194	536
Number of Unique Problems (Steps for Cognitive Tutor)	38,490	20,797	495,068
Number of Unique Users	27,760	875,492	3,292
Number of Responses	2,602,777	47,794,008	6,263,006
Average Number of Words per Problem	32	29	N/A
Information Types	Content+Context	Content+Context	Context
Granularity	Coarse-grained	Medium	Fine-grained

Table 2: Annotation results

Platforms	Fleiss' Kappa	Number of Skill Equivalencies
Khan Academy, ASSISTments	0.69	148
Khan Academy, Cognitive Tutor	0.72	85
ASSISTments, Cognitive Tutor	0.83	222

Preprocessing. For all three datasets, we drop the rows without skill assignment. For problem texts we tokenize words and clean the texts by removing stop words and converting to lower case. The basic descriptives of the datasets after preprocessing are shown in Table 1.

Skill Equivalence Labeling. There are currently no skill equivalency labels across these three platforms, so we create the labels ourselves. For each pair of platforms, the three authors separately annotated equivalent skills across platforms and then agreement was measured by Fleiss' Kappa and any conflict was resolved by discussion and majority voting. To give an example, the authors initially disagreed on whether skill "Angles 2" in Khan Academy should be mapped to skill "Angles - Obtuse, Acute, and Right" in ASSISTments. After discussion and inspecting the associated problems, a consensus was arrived at that they should not be linked as the Khan Academy skill involves angle calculation which is not needed in the ASSISTments skill. Since the granularity of taxonomies varies, there is likely to be one-to-many and many-to-one relationships between taxonomies. To account for such asymmetric relationships, we calculate Fleiss' Kappa by considering every possible pair of skills as "subject", and a binary value indicating equivalent or not as label. The Fleiss' Kappas and number of equivalent skill pairs are displayed in Table 2. The high Kappas suggest a good agreement among the authors.

4.2 Evaluation Metrics

We evaluate our models using two metrics, $\text{recall}@k$ and mean reciprocal rank. The skill taxonomies in our platform datasets differ in granularity, resulting in some skills having a one-to-many relationship across taxonomies. Thus, in addition to using a somewhat standard metric of $\text{recall}@k$ [23], we also report mean reciprocal rank to indicate the rank of the first true positive result.

$\text{Recall}@k$, where k is a user definable integer, measures the percentage of relevant destination skills that are contained within the model's top k predictions for each source skill, as calculated by $\frac{\text{true positives}@k}{\text{true positives}@k + \text{false negatives}@k}$. In our experiments, we choose k to be 5, which will return the 5 most similar predicted skills.

Mean reciprocal rank (MRR) calculates the rank of the first relevant destination skill for each source skill and takes the mean of the reciprocals of these ranks as output. It is given by $\frac{1}{|S|} \sum_{s \in S} \frac{1}{\text{rank}(s)}$ where S is the set of all source skills and $\text{rank}(s)$ denotes the rank of the first relevant destination skill for source skill s .

4.3 Experiment Details

4.3.1 Validation and Testing. We take different validation and testing strategies for different representation models, depending on whether they require hyperparameter tuning and whether they need ground truth labels to train. We apply hyperparameter search on validation data separate from the test set.

For Bag-of-words, TF-IDF, and Content2vec with pretrained word vectors, no hyperparameters or ground truth labels are involved in training. Test set results are produced by comparing the predictions against all the ground truth pairs.

For Content2vec with our own word vectors trained on problem texts, no skill pair labels are needed, but the word vector training process involves hyperparameters. Using 10-fold cross-validation, we search the best hyperparameters on the training set of each phase of the cross-validation and evaluate on the phase's test set with those hyperparameters, hence we get 10 best hyperparameter sets for the 10 folds instead of a single best hyperparameter set.

For Skill2vec and TAMF, the training requires both hyperparameter tuning and ground truth skill pairs. We similarly conduct a nested cross-validation on the training data within each phase of the 10-fold cross-validation to tune hyperparameters. The outer cross-validation tier has the same splits with Content2vec in order to compare models.

For Content2vec+Skill2vec, we simply concatenate the tuned content vector and context vector with the best hyperparameters for each fold.

At the dataset level, we treat Khan Academy and ASSISTments as training platforms and run experiments on them to compare models. We then pick the best model and conduct a final test on the mappings to and from Cognitive Tutor. We choose Cognitive

Table 3: Skill equivalency prediction results using symmetric data sources and skill representations

Representation	K2A	K2A	A2K	A2K
	Recall@5	MRR	Recall@5	MRR
Bag-of-words	0.55	0.43	0.38	0.46
TF-IDF	0.51	0.44	0.43	0.51
Content2vec (pretrained)	0.33	0.21	0.14	0.25
Content2vec	0.64	0.49	0.44	0.55
Skill2vec	0.61	0.51	0.11	0.14
Content2vec+Skill2vec	0.80	0.63	0.28	0.35
TAMF	0.74	0.67	0.22	0.21

Tutor as the test platform because it has only context information, whereas ASSISTments and Khan Academy have both content and context with which we can do holistic experiments on various skill representations, including symmetric and asymmetric source and destination data sources. A test on Cognitive Tutor without any model tuning will validate (or invalidate) our conclusions drawn from the other two platforms.

4.3.2 Training. We use Adam optimizer with learning rate 0.001 to train the translation model. The loss to minimize is the mean of cosine distances of training ground truth skill pairs. The maximum number of epochs is 1,000. We also randomly take 20% out of the training data as validation. If the validation loss has not decreased in 100 epochs, we stop the training and keep the model with the smallest validation loss. The models are trained on one NVIDIA Maxwell architecture GPU and each training job takes around 10 seconds to finish, with a whole nested cross-validation with 24 hyperparameter sets to tune totaling 5 hours (10 seconds $\times$ 10 $\times$ 8 $\times$ 24 = 19200 seconds) to run.

5 RESULTS

This section presents the results of several experiments. Section 5.1 compares equivalent skill prediction performance of the six skill representations in Khan Academy and ASSISTments assuming the same source and destination information types. Section 5.2 examines situations where source and destination platforms have different types of skill information (e.g., content in one and context in the other). Section 5.3 evaluates skill prediction performance on Cognitive Tutor as a test set using the best models as observed from the Khan Academy and ASSISTments results. Section 5.4 introduces the ablation study results for each method and the performance variations with different hyperparameters.

5.1 Symmetric Data Sources

In this section, we assume the source and destination platforms have the same type of data available (content, context, or both), and compare the six skill representations detailed in Section 3 on the skill equivalency prediction tasks between Khan Academy and ASSISTments. The results are shown in Table 3, where "K2A" refers to mapping from Khan Academy to ASSISTments taxonomy, and "A2K" is the other direction. The main observations are as follows:

- The highest MRRs for both directions are above 0.5, meaning the true predictions are at the second place or higher on average.
- Among all content-based methods, Content2vec with word vectors trained on problem texts is consistently the best. Content2vec with pretrained word vectors perform poorly, perhaps due to the insufficiency of the pretrained corpus in capturing information on math problem texts.
- In the mapping direction from Khan Academy to ASSISTments, incorporating both content and context information in the representation is better than either alone. The best method depends on the evaluation metric: Content2vec+Skill2vec is preferred with recall@5, and TAMF is favorable with MRR.
- In the mapping direction from ASSISTments to Khan Academy, the best method is Content2vec. Moreover, sequence information is counterproductive since Content2vec+Skill2vec is worse than Content2vec.
- The direction from Khan Academy to ASSISTments has better results than the other direction. This is likely because Khan Academy has a finer-grained taxonomy than ASSISTments, which will be discussed further in Section 6.3.

5.2 Asymmetric Data Sources

In certain scenarios, the combination of content and context information for a digital learning platform will not be available. For example, Cognitive Tutor provides only context information publicly, and other platforms like Junyi Academy⁶ have only content information made available. While the platforms themselves likely store sequence data, a new platform with little to no sequence data may still wish to map to other taxonomies. Our framework can accommodate such circumstances: it allows different types of source and destination skill representations to be presented to the model. The machine translation will learn transformations between the two vector spaces provided that there are ground truth skill pairs to train on. Experiments are conducted to test how skill equivalency prediction performance differs when data type availability is not symmetric between source and destination. There are three types of information input for a given platform: content, context, and both, resulting in 6 different combinations of asymmetric source and destination representation types. Table 4 gives the results of the different representation scenarios. The best models in terms of MRR for each input type (Content2vec, Skill2vec, and TAMF) are used.

Comparing the results with Table 3, we see that the highest recall@5 and MRR with asymmetric data sources are lower than symmetric cases, so it's better to first consider symmetric sources if all information is available. However, asymmetric sources are still a viable option. On one hand, asymmetric sources can not be avoided if one platform has only content and the other has only context. Moreover, even if one platform has both types of information and the other has only content or context, where training with symmetric sources is an option, it's still worth trying to incorporate both information in the first platform. For example, as shown in Table 3 and 4, TAMF to Skill2vec (asymmetric) is better than Skill2vec to Skill2vec (symmetric).

⁶<https://github.com/junyiacademy/junyiexercise>

Table 4: Skill equivalency prediction results using asymmetric data sources and skill representations

Source Content	Source Context	Source Model	Destination Content	Destination Context	Destination Model	K2A Recall@5	K2A MRR	A2K Recall@5	A2K MRR
✓		Content2vec		✓	Skill2vec	0.71	0.57	0.07	0.13
✓		Content2vec	✓	✓	TAMF	0.71	0.55	0.19	0.18
	✓	Skill2vec	✓		Content2vec	0.47	0.35	0.11	0.17
	✓	Skill2vec	✓	✓	TAMF	0.69	0.54	0.15	0.14
✓	✓	TAMF	✓		Content2vec	0.69	0.60	0.24	0.22
✓	✓	TAMF		✓	Skill2vec	0.70	0.63	0.12	0.16

5.3 Testing on Cognitive Tutor

We choose the models for mapping between Cognitive Tutor and the other two platforms anticipated to be the best based on the results of mapping shown in Tables 3 and 4. For each mapping direction, given that Cognitive Tutor has only context data, we look for the best representation for the other platform. For example, when selecting for Cognitive Tutor to Khan Academy, we look at the results where Khan Academy is the destination and the source is context only, and find the content-only destination has the highest MRR. We therefore choose Content2vec for Khan Academy.

The skill equivalency prediction results are displayed in Table 5. Mapping from Cognitive Tutor to the other two platforms reach relatively high recall@5 and MRR, supporting our methodology and conclusions from experiments between Khan Academy and ASSISTments; however, the opposite directions are worse. We believe this is because Cognitive Tutor has the finest-grained skill taxonomy and nuanced information is lost when forcing a mapping from coarse-grained to fine-grained (i.e., one-to-many).

5.4 Hyperparameter Tuning

Hyperparameters are tuned in the experiments with symmetric data sources. As detailed in Section 4.3.1, three methods need hyperparameter tuning: Content2vec, Skill2vec, and TAMF. Although we do not find a single best hyperparameter set as a consequence of cross-validation, we can compare the average performance of each hyperparameter set on the validation set. The metric is MRR in hyperparameter tuning thanks to its insensitivity to the one-to-many relationship.

For Content2vec, the hyperparameters include vector dimension, window size, and minimum count. Surprisingly, the best hyperparameters for each mapping direction are the same across all 10 folds, which are vector dimension 100, window size 20, and minimum count 50 from Khan Academy to ASSISTments, and vector dimension 100, window size 20, and minimum count 30 from ASSISTments to Khan Academy. Generally speaking, larger vector dimension and larger window size are better, while minimum token count does not have much impact on the performance.

For Skill2vec, the source and destination hyperparameters can be different, but there is no minimum count as we need to keep all skills, so the hyperparameters are source vector dimension, source window size, destination vector dimension, and destination window size. Figure 3 shows how they affect validation MRR across all folds. We can see that vector dimension influences the results more significantly than window size, but there is no single rule to

choose vector dimension, since sometimes large vector dimension is favorable while sometimes small dimension is better.

Figure 3: Skill2vec performance with different hyperparameters

TAMF has two hyperparameters for each platform, k for half vector dimension and λ for regularization coefficient. Shown in Figure 4 is the performance with different hyperparameter sets. Larger vector dimension is preferred, and regularization coefficient does not greatly affect the overall skill equivalency prediction results.

Figure 4: TAMF performance with different hyperparameters

6 ANALYSIS

This section delves into inspecting the taxonomy translation and edge cases. Section 6.1 visualizes skills from all taxonomies mapped to a single shared space; Section 6.2 considers skills with no matching skills in the destination platform; and Section 6.3 summarizes the effect of taxonomy granularity on skill equivalency prediction.

6.1 Visualization of Mapped Skills

We map skills from all platform taxonomies to a common space and visualize them to help understand skill equivalency in our taxonomy mapping models. Specifically, we project Khan Academy TAMF vectors and Cognitive Tutor Skill2vec vectors onto the

Table 5: Skill equivalency prediction results on Cognitive Tutor

Source	Source Representation	Destination	Destination Representation	Recall@5	MRR
Cognitive Tutor	Skill2vec	Khan Academy	Content2vec	0.72	0.72
Khan Academy	TAMF	Cognitive Tutor	Skill2vec	0.23	0.39
Cognitive Tutor	Skill2vec	ASSISTments	TAMF	0.88	0.79
ASSISTments	TAMF	Cognitive Tutor	Skill2vec	0.03	0.12

ASSISTments TAMF vector space and keep ASSISTments TAMF vectors unchanged, since the skill mapping in this direction gives the best equivalency prediction performance. In this space, we run k-means clustering to explore relationships among closely mapped skills. The number of clusters is set to 20 as determined by the "elbow method" heuristic. Finally, we apply t-SNE to reduce dimensionality and display the results in Figure 5. Skills are colored based on k-means applied to the original, high dimensional skill vectors.

We rank these clusters by a heuristic score: the percent of skills whose true matching skills from another taxonomy are also in the same cluster. The best cluster (score 1.0), the worst cluster (0.36), and a middle cluster (0.71) are featured in Figure 5. The best cluster is tightly grouped, containing only unit conversion skills. The worst cluster is the most dispersed, yet it still groups similar skills together like inequality and number line. The middle cluster is mostly about Pythagorean Theorem, with a few outliers. In this cluster, we can observe the fined-grained step-level skills of Pythagorean Theorem like calculating lengths and squares of hypotenuse and legs from Cognitive Tutor are clustered together with the coarse-grained problem-level skills from the other two platforms. The clustering results indicate that our models do map similar skills close to each other as desired. The complete visualization with cluster assignments and plot inspection tool can be found online⁷.

6.2 Skills without Counterparts

Owing to the uniqueness of each learning platform, some skills in one platform may not have a matching skill in another. In our case, 30% of skills in Khan Academy have no corresponding skills in ASSISTments and 31% of skills in ASSISTments have no counterparts in Khan Academy. We call these skills the "None" skills and they were ignored in the previous evaluations. While an algorithm can always return the most similar but not sufficiently equivalent skills, it is desirable to find a way for our method to report that there is no reasonable match. Therefore, we also conduct experiments to distinguish these untranslatable skills. Our method is to simply add a similarity threshold, whereby any prediction below that threshold will be considered as a "None" skill prediction, and the ground truth for those untranslatable skills is a mapping to the "None" skill. The evaluation metric used is again recall@5. We also compare this method against a baseline "Ignore", which is to train and predict as if all skills are translatable but consider the predictions for those "None" skills as wrongly predicted in evaluation.

Table 6 shows the results for Content2vec, Skill2vec, and Content2vec+Skill2vec, which are the best method with recall@5 for each input type. Considering "None" skills leads to a decrease in

Table 6: Skill equivalency prediction results including "None" skills

Representation	"None" Skill Strategy	K2A Recall@5	A2K Recall@5
Content2vec	Ignore	0.45	0.30
Content2vec	Threshold	0.53	0.42
Skill2vec	Ignore	0.43	0.08
Skill2vec	Threshold	0.47	0.13
Content2vec+Skill2vec	Ignore	0.56	0.24
Content2vec+Skill2vec	Threshold	0.62	0.26

performance compared to previously shown evaluations, underscoring the difficulty of addressing this issue. The threshold approach, however, performs slightly better than the baseline of not modeling None skills and counting them wrong. Given the simplicity of the threshold method, we hope future research can improve on this approach to classifying untranslatable skills between platforms.

6.3 Impact of Taxonomy Granularity on Skill Equivalency Prediction

We observed that there is a large performance discrepancy between directions of taxonomy mapping. For example, the best MRR from Khan Academy to ASSISTments is 0.67 achieved with TAMF, but the best MRR in the other direction is only 0.55. This phenomenon is also observed in other platform pairs (Table 7). For example, results of Cognitive Tutor to Khan Academy or ASSISTments are better than Khan Academy or ASSISTments to Cognitive Tutor.

A common characteristic of the poorer performing pairs is when mapping in the direction of a fine-grained taxonomy to a coarse-grained one. Cognitive Tutor is the finest with skills assigned to step, Khan Academy is the middle, and ASSISTments uses the coarsest-grained taxonomy of the three. We can observe the magnitude of the difference in performance between directions differs, as shown in Table 7. Cognitive Tutor and ASSISTments have larger difference than Cognitive Tutor and Khan Academy, and Khan Academy and ASSISTments. This result suggests that the more distant two platforms are in granularity, the larger performance discrepancy there will be between the two mapping directions.

7 LIMITATIONS AND FUTURE WORK

There were a few limitations to this research that we believe can be improved in future studies. First, the ground truth cross-platform skills were labeled by the authors, not domain experts. Second, the taxonomies used are restricted within the mathematics domain.

⁷ <https://cahlr.github.io/skill-equivalency-visualization>

Figure 5: t-SNE visualization of skills from all three platform taxonomies projected onto the ASSISTments TAMF vector space

Table 7: Discrepancy of skill equivalency prediction results (MRR) between opposite directions

Fine-grained Platform	Coarse-grained Platform	Fine to Coarse	Coarse to Fine	Difference
Khan Academy	ASSISTments	0.67	0.55	0.12
Cognitive Tutor	Khan Academy	0.72	0.39	0.33
Cognitive Tutor	ASSISTments	0.79	0.12	0.67

This was in part because the majority of publicly available digital learning platform clickstream and content datasets are from mathematics platforms. Third, we only utilized text content in our content-based approaches. Utilizing images, as was done in [5],

would capture a more complete representation of problems. Finally, pretrained contextual word embedding models like BERT [8] might be effective in improving content-based representations of problems, thus boosting skill equivalency prediction performance.

8 CONCLUSION

In this research, we demonstrated the viability of learning skill equivalencies across several taxonomies using data from the content of the problems skills are associated with and the clickstream sequences (i.e., contexts) in which those problems appear on a digital learning platform. We represented skills as vectors and employed machine translation to map between skill vector spaces, and validated the methodology on three digital learning platforms: ASSISTments, Khan Academy, and Cognitive Tutor.

We found that attempting to map from a coarser-grained taxonomy to a finer-grained taxonomy was considerably more difficult, with the best ASSISTments to Khan Academy recall@5 markedly lower than the reverse direction (0.44 vs 0.80) and, similarly, when mapping to the finer-grained taxonomy of the Cognitive Tutor (0.23 and 0.03 with Khan Academy and ASSISTments as sources, respectively). We also found that skill equivalence prediction was more accurate in experiments where there was symmetric data used to represent skills in both the source and destination taxonomies; however, skill prediction with asymmetric data performed comparably when mapping from fine-grained to coarse-grained taxonomies, particularly in the case of Cognitive Tutor (using Skill2vec) to ASSISTments (using TAMF), which scored 0.88 in recall@5.

These results are promising, but do not yet reach the level required for completely unattended automatic taxonomy mapping. The approach can, however, effectively triage a skill mapping or taxonomy crosswalk process and likely reduce the manual labor needed by a considerable amount. Use of the methods introduced in this paper, amplified by improvements defined in future work, could facilitate a more connected digital learning ecosystem, providing greater acknowledgement of students' prior learning and subsequently, more effective pedagogy.

REFERENCES

- [1] National Governors Association et al. 2010. Common Core State Standards. Washington, DC (2010).
- [2] Dmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- [3] Norman Bier, Sean Lip, Ross Strader, Candace Thille, and Dawn Zimmaro. 2014. An Approach to Knowledge Component/Skill Modeling in Online Courses. *Open Learning* (2014), 1–14.
- [4] Hao Cen, Kenneth R Koedinger, and Brian Junker. 2007. Is Over Practice Necessary?—Improving Learning Efficiency with the Cognitive Tutor through Educational Data Mining. In *Proceedings of the 2007 Conference on Artificial Intelligence in Education*. 511–518.
- [5] Devendra Singh Chaplot, Christopher MacLellan, Ruslan Salakhutdinov, and Kenneth Koedinger. 2018. Learning Cognitive Models Using Neural Networks. In *International Conference on Artificial Intelligence in Education*. Springer, 43–56.
- [6] Namyoun Choi, Il-Yeol Song, and Yongjun Zhu. 2016. A Model-Based Method for Information Alignment: A Case Study on Educational Standards. *Journal of Computing Science and Engineering* 10, 3 (2016), 85–94.
- [7] David T Conley. 2011. Crosswalk Analysis of Deeper Learning Skills to Common Core State Standards. *Educational Policy Improvement Center (NJ1)* (2011).
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT, Volume 1 (Long and Short Papers)*. 4171–4186.
- [9] Cambridge Assessment International Education. 2018. The 2018 Cambridge International Global Education Census. (2018).
- [10] Stephen E Fancsali, Michael V Yudelson, Susan R Berman, and Steven Ritter. 2018. Intelligent Instructional Hand Offs. *International Educational Data Mining Society* (2018).
- [11] Christian Fischer, Zachary A Pardos, Ryan Shaun Baker, Joseph Jay Williams, Padhraic Smyth, Renzhe Yu, Stefan Slater, Rachel Baker, and Mark Warschauer. 2020. Mining Big Data in Education: Affordances and Challenges. *Review of Research in Education* 44, 1 (2020), 130–160.
- [12] Neil T Heffernan and Cristina Lindquist Heffernan. 2014. The ASSISTments Ecosystem: Building a Platform that Brings Scientists and Teachers together for Minimally Invasive Research on Human Learning and Teaching. *International Journal of Artificial Intelligence in Education* 24, 4 (2014), 470–497.
- [13] Weijie Jiang and Zachary A Pardos. 2020. Evaluating Sources of Course Information and Models of Representation on a Variety of Institutional Prediction Tasks. *International Educational Data Mining Society* (2020).
- [14] Kenneth R Koedinger, Albert T Corbett, and Charles Perfetti. 2012. The Knowledge-Learning-Instruction Framework: Bridging the Science-Practice Chasm to Enhance Robust Student Learning. *Cognitive Science* 36, 5 (2012), 757–798.
- [15] Omer Levy and Yoav Goldberg. 2014. Neural Word Embedding as Implicit Matrix Factorization. In *Advances in Neural Information Processing Systems* 27. Curran Associates, Inc., 2177–2185.
- [16] Joshua J Michalenko, Andrew S Lan, and Richard G Baraniuk. 2017. Data-Mining Textual Responses to Uncover Misconception Patterns. In *Proceedings of the Fourth (2017) ACM Conference on Learning@Scale*. 245–248.
- [17] Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013. Exploiting Similarities among Languages for Machine Translation. *arXiv preprint arXiv:1309.4168* (2013).
- [18] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Advances in Neural Information Processing Systems*. 3111–3119.
- [19] Sara Morsy and George Karypis. 2019. Will This Course Increase or Decrease Your GPA? Towards Grade-Aware Course Recommendation. *Journal of Educational Data Mining* 11, 2 (2019), 20–46.
- [20] Liangming Pan, Xiaochen Wang, Chengjiang Li, Juanzi Li, and Jie Tang. 2017. Course Concept Extraction in MOOCs via Embedding-based Graph Propagation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 875–884.
- [21] John F Pane, Beth Ann Griffin, Daniel F McCaffrey, and Rita Karam. 2014. Effectiveness of Cognitive Tutor Algebra I at Scale. *Educational Evaluation and Policy Analysis* 36, 2 (2014), 127–144.
- [22] Zachary A Pardos, Yoav Bergner, Daniel T Seaton, and David E Pritchard. 2013. Adapting Bayesian Knowledge Tracing to a Massive Open Online Course in edX. In *6th International Conference on Educational Data Mining, EDM 2013*. International Educational Data Mining Society.
- [23] Zachary A Pardos, Hung Chau, and Haocheng Zhao. 2019. Data-Assistive Course-to-Course Articulation Using Machine Translation. In *Proceedings of the Sixth (2019) ACM Conference on Learning@Scale*. 1–10.
- [24] Zachary A Pardos and Anant Dadu. 2017. Imputing KCs with Representations of Problem Content and Context. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*. 148–155.
- [25] Zachary A Pardos, Zihao Fan, and Weijie Jiang. 2019. Connectionist Recommendation in the Wild: On the Utility and Scrutability of Neural Networks for Personalized Course Guidance. *User Modeling and User-Adapted Interaction* 29, 2 (2019), 487–525.
- [26] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online Learning of Social Representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 701–710.
- [27] Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. 2015. Deep Knowledge Tracing. In *Advances in Neural Information Processing Systems*. 505–513.
- [28] Darrell Porcello and Sherry Hsi. 2013. Crowdsourcing and Curating Online Education Resources. *Science* 341, 6143 (2013), 240–241.
- [29] Leena Razzaq, Neil T Heffernan, Mingyu Feng, and Zachary A Pardos. 2007. Developing Fine-Grained Transfer Models in the ASSISTment System. *Technology, Instruction, Cognition & Learning* 5, 3 (2007).
- [30] John Stamper and Zachary A Pardos. 2016. The 2010 KDD Cup Competition Dataset: Engaging the Machine Learning Community in Predictive Learning Analytics. *Journal of Learning Analytics* 3, 2 (2016), 312–316.
- [31] Cathlyn Stone, Abigail Quirk, Margo Gardener, Stephen Hutt, Angela L Duckworth, and Sidney K D'Mello. 2019. Language as Thought: Using Natural Language Processing to Model Noncognitive Traits that Predict College Success. In *Proceedings of the 9th International Conference on Learning Analytics & Knowledge*. 320–329.
- [32] Mega Subramaniam, June Ahn, Amanda Waugh, Natalie Greene Taylor, Allison Druin, Kenneth R Fleischmann, and Greg Walsh. 2013. Crosswalk between the "Framework for K-12 Science Education" and "Standards for the 21st-Century Learner": School Librarians as the Crucial Link. *School Library Research* 16 (2013).
- [33] Mark Wilson. 2004. *Constructing Measures: An Item Response Modeling Approach*. Routledge.
- [34] Cheng Yang, Zhiyuan Liu, Deli Zhao, Maosong Sun, and Edward Y Chang. 2015. Network Representation Learning with Rich Text Information. In *Proceedings of the 24th International Conference on Artificial Intelligence*. 2111–2117.
- [35] Ozgur Yilmazel, Niranjana Balasubramanian, Sarah C Harwell, Jennifer Bailey, Anne R Diekema, and Elizabeth D Liddy. 2007. Text Categorization for Aligning Educational Standards. In *2007 40th Annual Hawaii International Conference on System Sciences (HICSS'07)*. IEEE, 73–73.