

Prescribing Deep Attentive Score Prediction Attracts Improved Student Engagement

Youngnam Lee^{*1}, Byungsoo Kim^{*1}, Dongmin Shin¹, JungHoon Kim¹,
Jineon Baek^{1,2}, Jinhwan Lee¹ Youngduck Choi^{1,3}

¹Riiid! AI Research, ²University of Michigan, ³Yale University

{yn.lee, byungsoo.kim, dm.shin, junghoon.kim, jineon.baek, jh.lee, youngduck.choi}@riiid.co

ABSTRACT

Intelligent Tutoring Systems (ITSs) have been developed to provide students with personalized learning experiences by adaptively generating learning paths optimized for each individual. Within the vast scope of ITS, score prediction stands out as an area of study that enables students to construct individually realistic goals based on their current position. Via the expected score provided by the ITS, a student can instantaneously compare one's expected score to one's actual score, which directly corresponds to the reliability that the ITS can instill. In other words, refining the precision of predicted scores strictly correlates to the level of confidence that a student may have with an ITS, which will evidently ensue improved student engagement. However, previous studies have solely concentrated on improving the performance of a prediction model, largely lacking focus on the benefits generated by its practical application. In this paper, we demonstrate that the accuracy of the score prediction model deployed in a real-world setting significantly impacts user engagement by providing empirical evidence. To that end, we apply a state-of-the-art deep attentive neural network-based score prediction model to *Santa*, a multi-platform English ITS with approximately 780K users in South Korea that exclusively focuses on the TOEIC (Test of English for International Communications) standardized examinations. We run a controlled A/B test on the ITS with two models, respectively based on collaborative filtering and deep attentive neural networks, to verify whether the more accurate model engenders any student engagement. The results conclude that the attentive model not only induces high student morale (e.g. higher diagnostic test completion ratio, number of questions answered, etc.) but also encourages active engagement (e.g. higher purchase rate, improved total profit, etc.) on *Santa*.

Keywords

Intelligent Tutoring System, Score Prediction, Engagement, Deep Learning, Transformer

Youngnam Lee, Byungsoo Kim, Dongmin Shin, Junghoon Kim, Jineon Baek, Jinhwan Lee and Youngduck Choi "Prescribing Deep Attentive Score Prediction Attracts Improved Student Engagement" 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. 743 - 748

1. INTRODUCTION

The significance that standardized examinations (e.g. SAT and TOEIC) currently hold is to provide an objective criteria in which each individual's academic performance is measured. Accordingly, Intelligent Tutoring Systems (ITSs), which generate optimized learning paths for each student, often include functions such as estimating expected performance on standardized examinations. In this regard, measuring the expected academic performance of a student has become an interesting area of study in Artificial Intelligence in Education (AIED). These studies focus on modelling a student's understanding of a target subject based on their learning activities. For instance, Matrix Factorization (MF) [10, 22, 16, 17, 23, 7, 25, 24] is a prevalent method used for grade prediction, in which the latent vectors of students and courses are learned by factorizing a student-grade matrix into two low-rank matrices. Markov and semi-Markov models are also some other popular approaches for grade prediction [11, 7, 23]. With the advances in deep learning, neural network based models with deeper hidden layers, such as Multi-Layer Perceptron, Recurrent Neural Networks and Convolutional Neural Networks, were introduced to predict student's academic performance [21, 9, 11, 8]. In [3], the Transformer-based [29] bidirectional encoder model was first pre-trained to predict masked assessments and then fine-tuned to predict exam score, resulting in a state-of-the-art score prediction model. Although precision of academic performance prediction is significant as it is directly associated to a reliability of an ITS, previous studies have mainly focused on improving the accuracy of the prediction, leaving discussion about the benefits of precise prediction on student engagement fairly opaque.

In this paper, we direct our attention towards the correlation between the precision of score prediction and student engagement. Our study starts by hypothesizing that students will show higher level of engagement if they experience a more precise score prediction while interacting with ITS. We empirically verify our hypothesis on *Santa*, a multi-platform English ITS with approximately 780K users in South Korea that exclusively focuses on the TOEIC (Test of English for International Communications) Listening and Reading Test Preparation. In the experimental studies, we run a controlled A/B test with two score prediction models that differ in accuracy, which are respectively based on collab-

*Equal contribution.

orative filtering with Mean Absolute Error (MAE) of 78.9 and deep attentive neural networks with MAE of 49.8. The results show that the superior performing, deep attentive neural network based score prediction model induces more student engagement. These benefits range from ones that are derived from learning behavior (e.g. preliminary test completion ratio, membership rates, the average number of questions a student answered after the diagnostic test) to more active engagement (e.g. purchase rate, average revenue per user, and total profit). To the best of our knowledge, this is the first work studying the benefits of accurate score prediction of ITS on student engagement.

2. RELATED WORKS

The related works of this study can be grouped into two categories: academic performance prediction and student engagement.

2.1 Academic Performance Prediction

Predicting a student's academic performance is a significant aspect in solving the problems within AIED. A successful prediction model can be used to recommend appropriate courses, provide interventions for at-risk students, and optimally allocate learning materials. Extensive work has been conducted on performance prediction, exploring a wide range of methodologies from simple regressions to deep learning.

The most widely used methodology in grade prediction is low rank Matrix Factorization (MF) [10, 22, 16, 17, 23, 7, 25, 24]. Low rank MF assumes that there is a low-dimensional latent space containing features that can effectively represent both students and the academic tasks students will be graded on. These features can be interpreted as representations of a student's knowledge. We find these features by decomposing a student-grade matrix into a product of two low-rank matrices. The authors of [22] show that the MF-based model outperforms other course/student-specific regression models. [16] improved the model by assuming that different courses share a common latent feature space, since the totality of a student's knowledge should not change based on the courses they are taking.

Markov and semi-Markov models are another popular set of models for grade prediction [11, 7, 23]. These models capture the dynamic evolution of a student's learning status and leverage it to effectively predict outcomes. [7] develops course-specific hidden Markov and semi-Markov models for grade prediction. [11] models student behavior in MOOCs by using Hidden Markov Models (HMMs) and Multinomial Mixture Models (MMMs) to cluster sequences of student actions. The study applies an LSTM model to predict the students' final grades. Markov models are also used to estimate a student's performance on educational games [28] or to predict student retention in MOOCs [1].

[21, 9, 11, 8] introduce deep-learning based prediction models. The authors of [9] introduce two types of Bayesian deep learning models for grade prediction using Multi-Layer Perceptron and LSTM architectures. Their results show that their model outperforms several baseline models (including MF-based models and course-specific regression models) in detecting at-risk students. The authors of [3] propose As-

essment Modeling (AM), a pre-training method applicable to general ITSs. In AM, a model is first pre-trained to predict several assessments of a student automatically made by ITS during one's learning process. Their results show that a Transformer [29] based neural network model with AM improves model accuracy compared to the same network with other state-of-the-art pre-training methods (such as BERT [5] based word embedding and QuesNet [31] question embedding) on exam score prediction and review correctness prediction.

2.2 Student Engagement

Student engagement is also an actively studied topic in the field of AIED. Several works have analyzed student engagement patterns to figure out which factors vastly impact engagement. [30] studied how people use digital textbooks and compare engagement patterns among high school students, college students, and online website viewers. [18] investigated student engagement in an online learning system which outperformed a traditional classroom on key indicators of engagement, such as time on-task, engaged concentration, and boredom. [26] found correlations between semantic features of mathematics problems and indicators of engagement. [14] discriminated behavioral engagement and cognitive engagement, and argued that most of students who were behaviorally engaged were not cognitively engaged.

Another line of student engagement research focused on predicting engagement level. [20] proposed a two-phased approach for automatic engagement detection, which utilized contextual logs and appearance information to infer behavioral engagement. [19] investigated the relationship between engagement and performance. Firstly, this work analyzed log traces for each learner to calculate engagement indicators that represent learner's engagement level. Based on the quantified engagement indicators, prediction on the learner's performance were attempted.

Enrollment is a sign of strong engagement since it involves determination that a student must invest to take a certain course. Accordingly, predicting and promoting enrollment is highly relevant to student engagement research. [6] proposed a novel extension of Factorization Machines to infer students' course enrollment information from incomplete data. [2] presented a course enrollment recommender system which recommended selective and optional courses based on students' skills, knowledge and interests. [27] identified factors that affect the likelihood of enrolling. This work analyzed the enrollment predictability of such factors using logistic regression, support vector machines, and semi-supervised probability methods.

With the development of Massive Open Online Courses (MOOCs), several works studied student engagement in a MOOCs environment. [12] proposed a recommender which provides each student with an individual list of contacts based on their own profile and activities to foster their engagement in MOOCs. [15] investigated the relationship between students' self-evaluation of their previous knowledge and students' engagement behaviors in MOOCs through a polytomous item response theory model.

3. SCORE PREDICTION MODELS

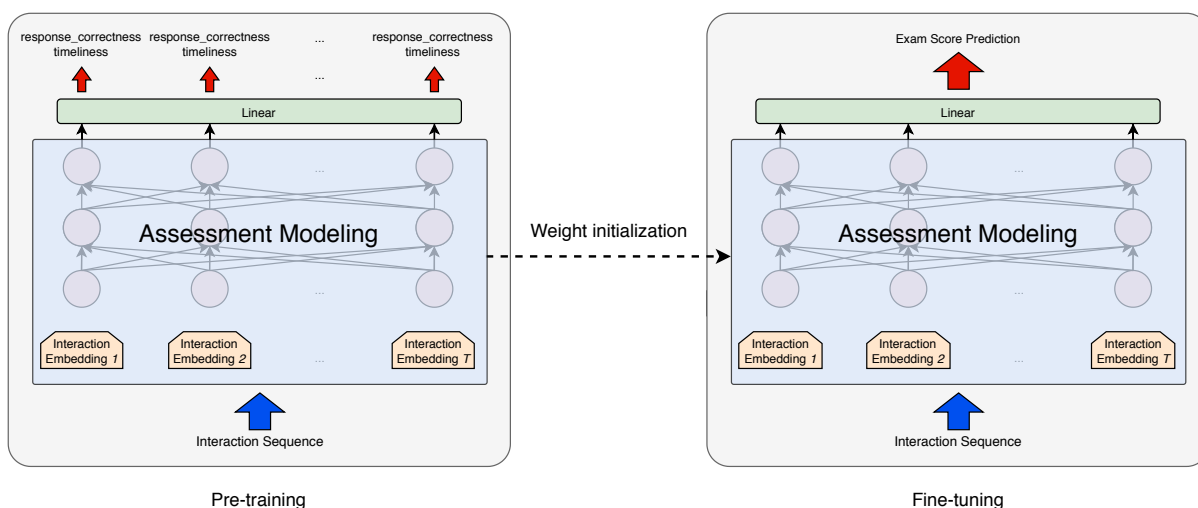


Figure 1: Pre-training/fine-tuning scheme of Assessment Modeling for score prediction. First, a model is pre-trained to predict two assessments: response correctness and timeliness. After pre-training, the last layer of the model is replaced with a layer with randomly initialized weights and appropriate dimension for score prediction. The parameters in the model are fine-tuned to predict exam scores.

Our studies are based on comparing the two approaches for score prediction: a collaborative filtering based approach and Assessment Modeling. The following subsections briefly cover each approach. More detailed descriptions can be found in [13] and [3].

3.1 Collaborative Filtering based Approach

There are two phases in the Collaborative Filtering (CF)-based score prediction approach. First, the CF-based model developed in [13] estimates the probability that a student responds correctly to each potential question. In this model, each user or question is represented as a k -dimensional latent vector, where k is the number of hidden concepts. For instance, if there are n users with m questions, we have user vectors L_1, L_2, \dots, L_n and question vectors R_1, R_2, \dots, R_m each with dimension k . The knowledge level of user i understanding question j is represented as $X_{ij} = L_i \cdot R_j$. Accordingly, the probability of user i getting question j correct is modeled as

$$\phi(X_{ij}) = \phi_a + \frac{1 - \phi_a}{1 + e^{-\phi_c(X_{ij} - \phi_b)}},$$

where ϕ_a , ϕ_b , and ϕ_c are parameters appropriately set, independently of questions or users. The learning algorithm then finds the maximum likelihood estimator by minimizing the negative of log-likelihood of observed user-question entries with Frobenius norm regularizer terms through the projected stochastic gradient descent.

Given the response correctness probabilities calculated from the CF-based model, scores for Listening Comprehension (LC) and Reading Comprehension (RC) are calculated through the following quadratic equations

$$\begin{aligned} score_{LC} &= \theta_2 x_{LC}^2 + \theta_1 x_{LC} + \theta_0 \\ score_{RC} &= \theta_5 x_{RC}^2 + \theta_4 x_{RC} + \theta_3, \end{aligned}$$

where x_{LC} and x_{RC} are each the average of predicted response correctness probability of potential questions in LC and RC, and θ s are properly set parameters. The final score is the sum of $score_{LC}$ and $score_{RC}$.

3.2 Assessment Modeling

[3] introduced Assessment Modeling (AM), a fundamental pre-training method for general class of ITSs. The motivation behind the works of AM is to deal with label-scarce problems in AIED. Score prediction is a typical example of such label-scarce educational problems since standardized exam scores are not obtainable within ITS. Collecting the exam scores involves student action taken outside ITS. The approach proposed in [3] is based on a pre-training/fine-tuning paradigm. In the pre-training phase, the Transformer-based [29] bidirectional encoder model is trained to predict randomly masked assessments, which are interactive educational features available in ITS. Examples of these assessments include response correctness (whether a student provides a correct response to a given question) and timeliness (Whether a student responds to each question within the time limit specified by domain experts). In the fine-tuning phase, the last layer of the pre-trained model is replaced with a randomly initialized layer with an appropriate dimension for a specific downstream task. Afterwards, the parameters in the model are updated to predict labels in the downstream task. In the experimental studies conducted on EdNet [4], AM outperformed pre-training methods that learn the contents of learning materials in several downstream tasks including score prediction. See Figure 1 for graphical description of AM.

4. EXPERIMENTS

4.1 Santa service

Santa is a multi-platform English ITS with approximately 780K users in South Korea that exclusively focuses on the

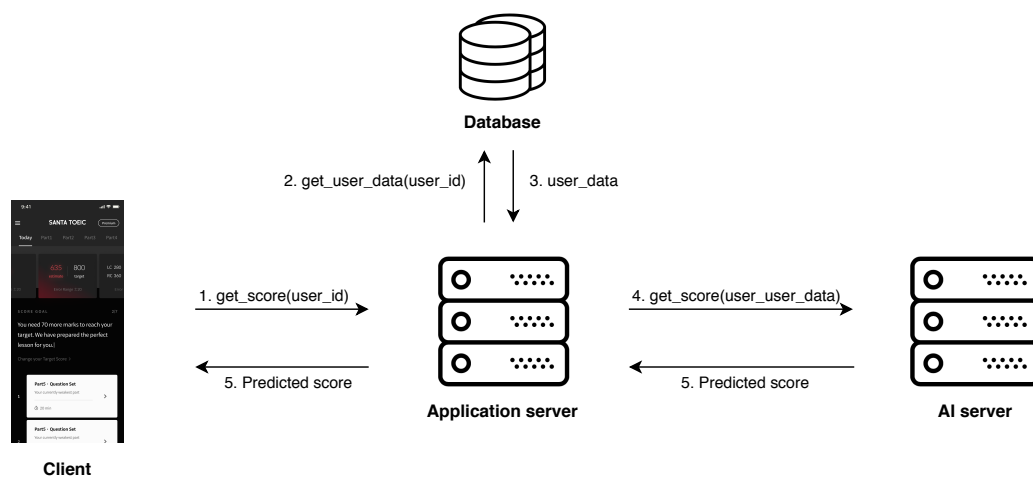


Figure 2: The flow of score prediction.

TOEIC (Test of English for International Communications) standardized examinations. TOEIC is an English proficiency test that consists of two timed sections (listening and reading) each with 100 questions that adds up to a combined total score between 0 to 990. Santa utilizes several AI techniques to optimize the preparation process of the TOEIC examination for students. When the application is first initiated, a preliminary placement test with 7 to 11 problems is given to diagnose the student’s current state and predict their expected score in real-time. After the diagnostic test, a user response prediction model is used to dynamically suggest problems which corresponds to the student’s current position within the TOEIC ladder. The prediction model is calculated by computing a user’s overall correctness rate, eliminating problems that students have answered correctly with high probability and then selecting the best possible content based on expert heuristics. Based on the user’s previous data, the predicted scores can be provided in various forms throughout the service, as shown in Figure 3. Figure 2 shows the flow of score prediction.

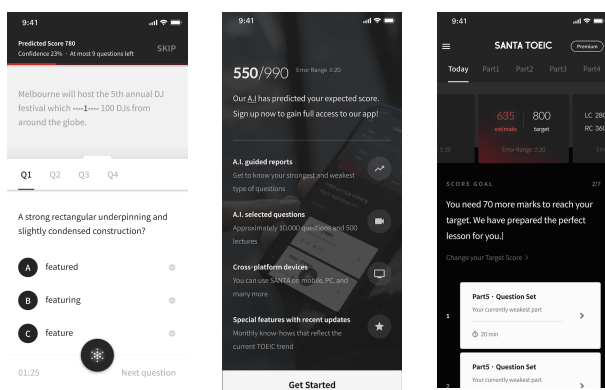


Figure 3: Various representations of predicted scores within the application.

4.2 Performance of Score Prediction Model

Santa has previously used a CF-based model for score prediction which has recently been replaced with a deep attentive model. To train the model, we aggregate the real TOEIC scores reported by users of Santa. Santa offered to reward to the users who have reported their score and was able to obtain a total of 2,594 score reports for 6 months. The data is then divided into a training set (1,302 users, 1815 labels), validation set (244 users, 260 labels), and a test set (466 users, 519 labels). We use EdNet as pre-training task data and the student sequence data as the label (TOEIC score). Table 4.2 shows the MAE (Mean Absolute Error) of the two models for the test set.

	CF	Deep Attentive model
MAE	78.91	49.84

Table 1: MAE of collaborative filtering and attentive model

4.3 A/B test setup

From February 24th to April 2nd, we conducted an A/B test by randomly administering two different score prediction algorithms to the application users: one based on a collaborative-filtering algorithm and another one based on deep-learning. 50,451 students were allocated to the collaborative filtering algorithm and 17,019 students were provided a deep-learning algorithm. We analyzed each student’s response and action (such as time of registration, question response time, purchase rate, etc.) to spot any noteworthy statistics that can validate our experiment.

4.4 Experimental Results

In this section, we discuss how a high quality of the predicted scores can significantly impact student morale.

4.4.1 Student Motivation

Our first test statistic is the preliminary test completion ratio. The completion rate of the initial placement test is a crucial indicator that could represent a student’s motivation,

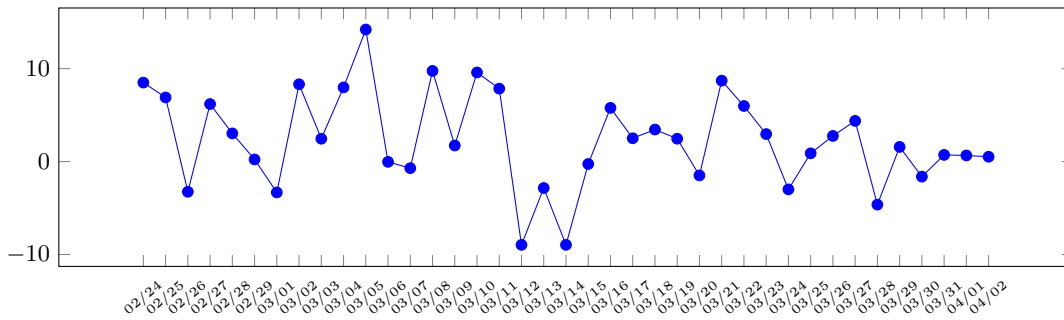


Figure 4: Comparison of the number of questions solved per day between the users of the A/B test.

as only students who are willing to learn will try to finish their diagnostic test. For each question a student answers in SANTA, a predicted score that is updated in real-time is projected on the top left corner. This allows for the user to immediately check the quality of the expected score, thus strengthening the trust that the user may have with the application. A/B test results show that the deep attentive model has a higher completion rate of 64.93% than the CF-model with 65.90%.

Next, we look at changes in membership rates. A membership rate of an application in a sense signifies greater magnitude of student motivation than the completion rate as it directly indicates the determination of a user who wishes to use the application. Out of a total of 67,470 users that have used Santa during the A/B test period, 44,297 users finished their diagnostic tests and 28,065 users have registered to sign up with the application. The A/B test shows that the deep attentive model has a registration rate of 43.13% while the CF-based model has 44.55%.

The average number of questions a user answered after the diagnostic test is also significant proof of a student’s educational drive. The A/B test results show that with a deep attentive model a student solved an average of 22.73 questions, while with a CF-based model the user only solved 20.03. Figure 4 shows the comparison of the number of questions answered per day between the users of the A/B test. The x-axis represents the date and the y-axis represents the gap between average number of questions answered in a deep attentive model and a CF-based model. If the gap is positive, the former model has on average more questions solved, and vice versa. We can observe that more questions from the deep attentive model were solved mostly throughout the A/B test time period.

	CF	Deep Attentive model
Completion rate (%)	64.93	65.90
Registration rate (%)	43.13	44.55
# of solved questions	20.03	22.73

Table 2: Experimental results of student motivation

4.4.2 Active Student Engagement

In this section, we demonstrate active student engagement based on different score prediction models via taking a look at the financial benefits the models bring. Monetary profits are an essential factor in evaluating a service, since it is

an important indicator of user engagement as a high level of user engagement directly results in financial success. We measure business impact with 3 metrics : purchase rate, Average Revenue Per User (ARPU), and total profit. In this context, purchase rate is defined as the number of users that decided to purchase full access to the app during the A/B test period. The test results show that the purchase rate for the deep attentive model was 2.73% while the CF-based model had a 2.37% rate, showing a 15.19% increase for the deep attentive model. For ARPU, the deep attentive model averaged \$3.23 whilst a CF-based model averaged \$2.83. Total profit during testing period also yielded \$162,933.88 for the former while it only gathered \$142,949.55 for the latter (since the two models had different parameters, these values were normalized based on the ratio of the model parameters). Comparing these 3 metrics, we conclude that the model with higher accuracy in the deep attentive model shows better results as well.

	CF	Deep Attentive model
Conversion rate (%)	2.37	2.73
ARPU (\$)	2.83	3.23
Total profit (\$)	142,949.55	162,933.88

Table 3: Experimental results of student engagement

5. CONCLUSIONS

Recent developments in ITS have enabled customized education by suggesting optimal strategies for individual students to approach studying. SANTA has also assisted its users to better prepare for the TOEIC English fluency standardized examinations by utilizing various learning techniques. Recently, SANTA has shifted from a collaborative-filtering model to a deep attentive model that has proved to be an upgrade over the former. To inquire about the benefits of using a fastidious model, this paper conducts various experiments and investigates their results. Analyzing the results of various experiments leads us to believe that deep attentive model entails a higher level of student motivation and engagement. Therefore, we claim that a more accurate model, in this case, the deep attentive model, could induce improved student engagement.

6. REFERENCES

- [1] G. Balakrishnan and D. Coetzee. Predicting student retention in massive open online courses using hidden markov models. *Electrical Engineering and Computer*

Sciences University of California at Berkeley, 53:57–58, 2013.

- [2] H. Bydžovská. Course enrollment recommender system. *International Educational Data Mining Society*, 2016.
- [3] Y. Choi, Y. Lee, J. Cho, J. Baek, D. Shin, S. Lee, Y. Cha, B. Kim, and J. Heo. Assessment modeling: Fundamental pre-training tasks for interactive educational systems, 2020.
- [4] Y. Choi, Y. Lee, D. Shin, J. Cho, S. Park, S. Lee, J. Baek, B. Kim, and Y. Jang. Ednet: A large-scale hierarchical dataset in education, 2019.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [6] J. P. González-Brenes and R. Edezhath. Inferring course enrollment from partial data. In *International Conference on Artificial Intelligence in Education*, pages 429–432. Springer, 2018.
- [7] Q. Hu and H. Rangwala. Course-specific markovian models for grade prediction. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 29–41. Springer, 2018.
- [8] Q. Hu and H. Rangwala. Academic performance estimation with attention-based graph convolutional networks. *arXiv preprint arXiv:2001.00632*, 2019.
- [9] Q. Hu and H. Rangwala. Reliable deep grade prediction with uncertainty estimation. *arXiv preprint arXiv:1902.10213*, 2019.
- [10] Z. Iqbal, J. Qadir, A. N. Mian, and F. Kamiran. Machine learning based student grade prediction: A case study. *arXiv preprint arXiv:1708.08744*, 2017.
- [11] Y. Jo, K. Maki, and G. Tomar. Time series analysis of clickstream logs from online courses. *arXiv preprint arXiv:1809.04177*, 2018.
- [12] H. Labarthe, F. Bouchet, R. Bachelet, and K. Yacef. Does a peer recommender foster students' engagement in moocs?. *International Educational Data Mining Society*, 2016.
- [13] K. Lee, J. Chung, Y. Cha, and C. Suh. Machine learning approaches for learning analytics: Collaborative filtering or regression with experts? In *NIPS Workshop, Dec*, pages 1–11, 2016.
- [14] Q. Li and R. Baker. Understanding engagement in moocs. In *EDM*, pages 605–606, 2016.
- [15] J. Lui and H. Li. Exploring the relationship between student pre-knowledge and engagement in moocs using polytomous irt. In *Proceedings of the 10th International Conference on Educational Data Mining*, pages 410–411. ERIC, 2017.
- [16] S. Morsy and G. Karypis. Cumulative knowledge-based regression models for next-term grade prediction. In *Proceedings of the 2017 SIAM International Conference on Data Mining*, pages 552–560. SIAM, 2017.
- [17] S. Morsy and G. Karypis. Sparse neural attentive knowledge-based models for grade prediction. *arXiv preprint arXiv:1904.11858*, 2019.
- [18] K. Mulqueeney, L. A. Mingle, V. Kostyuk, R. S. Baker, and J. Ocumpaugh. Improving engagement in an e-learning environment. In *International Conference on Artificial Intelligence in Education*, pages 730–733. Springer, 2015.
- [19] V. Naik and V. Kamat. Analyzing engagement in an on-line session. In *International Conference on Artificial Intelligence in Education*, pages 359–364. Springer, 2019.
- [20] E. Okur, N. Alyuz, S. Aslan, U. Genc, C. Tanriover, and A. A. Esme. Behavioral engagement detection of students in the wild. In *International Conference on Artificial Intelligence in Education*, pages 250–261. Springer, 2017.
- [21] A. P. Patil, K. Ganesan, and A. Kanavalli. Effective deep learning model to predict student grade point averages. In *2017 IEEE International Conference on Computational Intelligence and Computing Research (ICIC)*, pages 1–6. IEEE, 2017.
- [22] A. Polyzou and G. Karypis. Grade prediction with course and student specific models. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 89–101. Springer, 2016.
- [23] L. Rechkoski, V. V. Ajanovski, and M. Mihova. Evaluation of grade prediction using model-based collaborative filtering methods. In *2018 IEEE Global Engineering Education Conference (EDUCON)*, pages 1096–1103. IEEE, 2018.
- [24] Z. Ren, X. Ning, A. Lan, and H. Rangwala. Grade prediction based on cumulative knowledge and co-taken courses. In *Proceedings of the 12th International Conference on Educational Data Mining (EDM)*. ERIC, 2019.
- [25] Z. Ren, X. Ning, and H. Rangwala. Ale: Additive latent effect models for grade prediction. In *Proceedings of the 2018 SIAM International Conference on Data Mining*, pages 477–485. SIAM, 2018.
- [26] S. Slater, R. Baker, J. Ocumpaugh, P. Inventado, P. Scupelli, and N. Heffernan. Semantic features of math problems: Relationships to student learning and engagement. *International Educational Data Mining Society*, 2016.
- [27] A. Slim, D. Hush, T. Ojah, and T. Babbitt. Predicting student enrollment based on student and college characteristics. *International Educational Data Mining Society*, 2018.
- [28] M. Tadayon and G. Pottie. Predicting student performance in an educational game using a hidden markov model. *arXiv preprint arXiv:1904.11857*, 2019.
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [30] J. Warner, J. Doorenbos, B. Miller, and P. J. Guo. How high school, college, and online students differentially engage with an interactive digital textbook. In *EDM*, pages 528–531, 2015.
- [31] Y. Yin, Q. Liu, Z. Huang, E. Chen, W. Tong, S. Wang, and Y. Su. Quesnet: A unified representation for heterogeneous test questions. *arXiv preprint arXiv:1905.10949*, 2019.