
A Learned Simulation Environment to Model Student Engagement and Retention in Automated Online Courses

N. Imstepf¹, S. Senn¹, A. Fortin², B. Russell², and C. Horn¹

¹ Zurich University of Applied Sciences, ICLS, Schloss 1, 8820 Wädenswil, Switzerland

² Kikodo Education Technology Sàrl, Av. Louis Ruchonnet, Lausanne 1003, Switzerland

Abstract

We developed a simulator to quantify the effect of exercise ordering on both student engagement and retention. Our approach combines the construction of neural network representations for users and exercises using a dynamic matrix factorization method. We further created a machine learning models of success and dropout prediction. As a result, our system is able to predict student engagement and retention based on a given sequence of exercises selected. This opens the door to the development of versatile reinforcement learning agents which can substitute the role of private tutoring in exam preparation.

Keywords: education technology, student modeling, deep reinforcement learning.

1 Introduction

Survey data reveals that 50% of US and UK households would like to hire a tutor, but cannot afford it [10]. This issue is especially acute in the context of high-stakes testing such as GCSE and SAT/ACT which can have a determining effect on life outcomes of students, while results in such tests are also strongly influenced by the ability to hire a quality private tutor [6, 11].

Furthermore, a major concern for sustained economic growth today is the shortage of skilled workers [3]. Online learning combined with new approaches in education technology may help solve this challenge by improving access to the highest quality independent learning materials and processes. Global AI usage in education is estimated to have an annual growth rate of 36.6 percent during 2022-2030 and is predicted to reach \$25.77 billion by 2030 [16].

Classical classroom teaching often suffers from large percentages of disengaged students for which the presented material is either too difficult or too easy [8]. This problem is exacerbated by the diversity of the existing knowledge of the class. Thus, the promise of personalized learning is to optimize the learning path for each student individually, and therefore maximize learning, retention, and ultimately exam performance for each student.

While past AI-based approaches have used models trained by supervised learning, we propose to use intelligent agents which are able to continuously learn throughout the complete student journey. Training such agents requires the development of a simulator (aka digital twin) for student actions, which is the topic of this paper. It forms the basis for a natural next step, the development of an autonomous digital tutor.

Our main contributions are:

- An end-to-end approach to optimizing student interactions in online learning;
- A novel dynamic matrix factorization method to learn neural representations of users and exercises from user data;




Figure 1: Overview of our proposed modeling chain.

- A demonstration of how Model-Based reinforcement learning can be enabled by learning the components of a simulator.

The benefits of this approach translate into a more engaging learning experience for the students, higher performance, and increased retention rates.

2 Method

Our aim is to build a simulator that allows us to quantify the effect of different exercise orders on student engagement and retention. This could then be used to understand and improve the individual learning path of each student. In particular, such an approach provides the basis for the development of autonomous, intelligent agents to be deployed in automatic online courses. Using these agents, one can dynamically optimize student engagement and retention as a function of past interactions. The corresponding modeling chain is illustrated in Fig. 1. It comprises four main steps, (1) data gathering and pre-processing, (2) learning of neural representations for users and exercises, (3) modeling of success and dropout probability (4) construction of the student interaction simulator. All steps are described in detail below.

2.1 Collection and processing of student interaction data

The user question submission data was based on the Python learn-to-code site www.kikodo.io. For each user, a time series of "events" was used, where each event is the submission of a single exercise, and the outcome of that submission. Although additional metadata were present in the original data, only the correctness or incorrectness of the submission was considered.

The exercises on this specific e-learning site are structured simply into "workbooks", which are themselves grouped by topic. Each workbook contains a number of questions in the range of 10-100. Prior to the preparation of this study all learners had questions within a workbook presented to them in a random order.

The users of the site have been anonymized in all reporting and analysis.

2.2 Dynamic matrix factorization

Developing a neural network model to select the best exercise for a given user requires representations of users and exercises. A common way to calculate these from user-exercise interaction data is collaborative filtering (CF) as used in recommender systems [17]. While standard user-based and item-based CF just take the user-item¹ matrix to derive representations, matrix factorization actively optimizes the representations of both users and items at the same time to minimize the average prediction error [7]. The predicted scores per exercise can be computed as $\hat{S} = UE$, where $\hat{S} \in \mathbb{R}^{n \times m}$ is the user-exercise score matrix, $U \in \mathbb{R}^{n \times l}$ contains the user's latent factors and $E \in \mathbb{R}^{l \times m}$ the exercise's latent factors, and n = number of users, m = number of exercises, l = number of latent factors. The predicted exercise score for user u and exercise e can then be

¹In our case the items are the exercises.

Table 1: Performance comparison for different success prediction model variants

Model type	Number of last exercises used	RMSE
Random Forest	3	0.238
SVM	3	0.246
XGBoost	3	0.243
Random Forest	10	0.227
SVM	10	0.237
XGBoost	10	0.239

computed as:

$$\hat{s}_{ue} = \sum_{f \in [1:l]} U_{uf} E_{fe} \tag{1}$$

This ensures that our representations for users and exercises are consistent with each other. The problem then consists in optimizing the objective function

$$\operatorname{argmax}_{U,E} \|S - \hat{S}\| + \lambda_1 \|U\| + \lambda_2 \|E\| \tag{2}$$

where the last two terms provide regularization, and λ_1 , and λ_2 are hyperparameters. To perform the optimization, we use the gradient descent algorithm in its convenient implementation in the deep learning library PyTorch[15].

In practice, however, where most users only complete a single course with a relatively small number of exercises contained in a limited time period, the standard MF approach described above faces the problem that for any two separate time periods the overlap of users and exercises is small. Which amounts to a kind of continuous cold start problem[9]. To tackle it we propose *dynamic matrix factorization*, which dynamically extends the user-exercise matrix each time a new user or exercise appears in the interaction stream (adding a new row or column with values corresponding to an average exercise or user, respectively). We perform a small number of gradient descent training steps after each new observed interaction. This has three important consequences:

- It allows us to derive representations for new users and new exercises already after a single recorded interaction, effectively solving the continuous cold start problem;
- Additional observations will help to gradually improve the representations over time;
- Information learned from past users-exercise interactions can continuously flow into the representations of new users and exercises.

Our experiments confirm that representations indeed improve over time by comparing the performance of downstream models (see Sect. 2.3) at different numbers of user interactions.

2.3 Success and dropout prediction

To model a user’s interaction when presented with a new exercise, we need to predict her/his performance on the exercise as well as her/his decision to drop out, i.e. not to finish the exercise, and quite the course.

We use machine learning to predict these outcomes from the historical data. As the input for the score prediction model, we concatenate the following into a single vector:

$$u_t, e_{t-n}, s_{t-n}, \dots, e_{t-1}, s_{t-1}, e_t \tag{3}$$

where u_t and e_t are the user and exercise representations learned in Sect. 2.2, and s_t is the performance score. The performance of different machine learning algorithms is shown in Tab. 1. The random forest model using the representations of the last $n = 10$ exercises as input was selected for downstream tasks.

As the input for the dropout prediction model, we concatenate the following into a single vector:

$$u_t, e_{t-n}, s_{t-n}, \dots, e_{t-1}, s_{t-1}, e_t, s_t \tag{4}$$

Up-sampling of the minority class was used to deal with class imbalance. The performance of the dropout prediction model is shown by the ROC curve in Fig. 2.

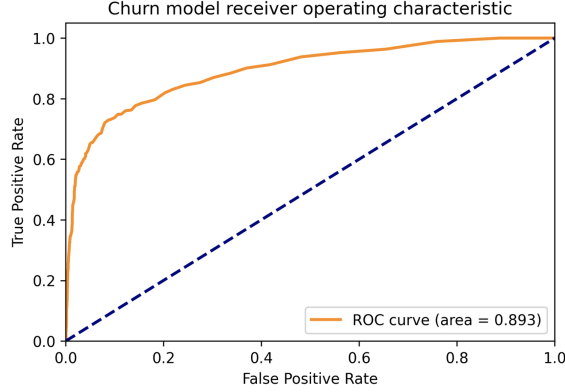


Figure 2: Performance of dropout prediction model, measured by the ROC curve[4].

2.4 Student interaction simulator

Given the models learned in Sect. 2.3 it is now possible to construct a student interaction simulator. It receives the information about the next exercise the student should solve as input and returns a dropout decision as well as a performance score for the given exercise. It also returns a reward which can be used to train a potential learning agent tasked with selecting the next exercise to present to the student. In order to guide the student towards a desired target performance score s_{target} and minimize the dropout probability, p_{dropout} , we define the reward function as:

$$R = \sum_e [(s_e - s_{\text{target}})^2 + \alpha(1 - p_{\text{dropout}})] \quad (5)$$

where s_e is the student’s performance score for exercise e . We compared the interaction with the simulation environment of two different strategies and observed the rewards they achieve

- Applying the same exercise order as in the historical data.
- Applying an AI-agent trained with PPO [21].

The results are shown in Fig. 3. We observe that the AI agent is able to choose orders of exercises with consistently higher reward per user. This shows a clear potential for AI agents trained with our approach to improve learning outcomes in automated online courses.

3 Related work

Education offers a large number of opportunities for applications of artificial intelligence [18]. The advancement of digitization and learning management systems produces an increasing amount of data [22]. Analyzing it is the subject of the field of educational data mining [20, 19]. E.g. it was found that dedicated feedback can enhance students’ motivation [2]. Machine learning (ML) has been applied in digital education in different ways, mostly using classical ML approaches [13]. A big, mostly untapped potential for reinforcement learning exists in e-learning [12]. Previous studies have mostly focused on student engagement [2], assuming that retention will follow. A number of studies focus on the prediction of student performance [1], as well as the prediction and reduction of student dropout [23]. The usefulness of recommendation systems for designing smart learning management system for digital learning has been shown in several studies, where the most popular methods are collaborative filtering and content-based filtering [14]. Recently, student profiling models have found increased attention, mostly using deep learning [5]. However, to date, no ML model exists which is able to actively steer the complete learning path of individual students in e-learning courses.

4 Conclusions

We developed a simulator for the automatic selection of exercises for individual students in automated online courses. We first use our dynamic matrix factorization method to derive neural representations




Figure 3: Comparison of accumulated scores of episode rewards over time achieved if exercises are chosen like given in the historical data (blue) and by the RL agent (orange). In each episode, a new user is randomly selected.

for users and exercises, which provide the input features for success and dropout prediction models, trained via supervised machine learning. The resulting models could successfully be used in a simulator which enables the training of reinforcement learning agents. Initial tests show potential for such agents to optimize the learning path for each individual student in automated online courses. Next, we plan to incorporate the optimization of feedback, coding hints and assessments into the automated learning path and test our AI agent in a real-world setting.

Acknowledgements

We thank Innosuisse for the support within 58831.1 INNO-ICT.

References

- [1] M. Anoopkumar and A. Rahman. A comprehensive survey on educational data mining and use of data mining techniques for improving teaching and predicting student performance. In *Education*, 2015.
- [2] Sarra Ayouni, Fahima Hajjej, Mohamed Maddeh, and Shaha Al-Otaibi. A new ML-based approach to enhance student engagement in online environment. *PLOS ONE*, 16(11):e0258788, November 2021.
- [3] James Bessen. Employers aren’t just whining—the “skills gap” is real. *Harvard Business Review*, 25, 2014.
- [4] Tom Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.
- [5] Chong Guan, Jian Mou, and Zhiying Jiang. Artificial intelligence innovation in education: A twenty-year data-driven historical analysis. *International Journal of Innovation Studies*, 4(4):134–147, December 2020.
- [6] Judith Ireson. Private tutoring: How prevalent and effective is it? *London Review of Education*, 2004.
- [7] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [8] June Kronholz. Can khan move the bell curve to the right? *Education Next*, 12, No. 2, 2012.

- [9] Blerina Lika, Kostas Kolomvatsos, and Stathes Hadjiefthymiades. Facing the cold start problem in recommender systems. *Expert Systems with Applications*, 41(4, Part 2):2065–2073, 2014.
- [10] Rebecca Montacute and Carl Cullinane. Parent power 2018: How parents use financial and cultural resources to boost their children’s chances of success. <https://www.suttontrust.com/wp-content/uploads/2019/12/Parent-Power-2018.pdf>, September 2018.
- [11] Raeal Moore, Edgar Sanchez, and Maria Ofelia San Pedro. Investigating test prep impact on score gains using quasi-experimental propensity score matching. act working paper 2018-6. *ACT, Inc.*, 2018.
- [12] Abdallah Moubayed, Mohammadnoor Injadat, Ali Bou Nassif, Hanan Lutfiyya, and Abdallah Shami. E-Learning: Challenges and Research Opportunities Using Machine Learning & Data Analytics. *IEEE Access*, 6:39117–39138, 2018.
- [13] Hussan Munir, Bahtijar Vogel, and Andreas Jacobsson. Artificial Intelligence and Machine Learning Approaches in Digital Education: A Systematic Revision. *Information*, 13(4):203, April 2022.
- [14] Dina Fitria Murad, Yaya Heryadi, Bambang Dwi Wijanarko, Sani Muhamad Isa, and Widodo Budiharto. Recommendation System for Smart LMS Using Machine Learning: A Literature Review. In *2018 International Conference on Computing, Engineering, and Design (ICCED)*, pages 113–118, Bangkok, Thailand, September 2018. IEEE.
- [15] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* 32, pages 8024–8035. Curran Associates, Inc., 2019.
- [16] psmarketresearch. Ai in education market research report. <https://www.psmarketresearch.com/market-analysis/ai-in-education-market>, December 2022.
- [17] Francesco Ricci, Lior Rokach, and Bracha Shapira. *Introduction to Recommender Systems Handbook*, pages 1–35. Springer US, Boston, MA, 2011.
- [18] Ido Roll and Ruth Wylie. Evolution and revolution in artificial intelligence in education. *International Journal of Artificial Intelligence in Education*, 26(2):582–599, June 2016.
- [19] Cristobal Romero and Sebastian Ventura. Data mining in education. *WIREs Data Mining and Knowledge Discovery*, 3(1):12–27, 2013.
- [20] Cristóbal Romero and Sebastián Ventura. Educational Data Mining: A Review of the State of the Art. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 40(6):601–618, November 2010.
- [21] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms. Technical Report arXiv:1707.06347, arXiv, August 2017. arXiv:1707.06347 [cs] type: article.
- [22] Filippo Sciarrone. Machine Learning and Learning Analytics: Integrating Data with Learning. In *2018 17th International Conference on Information Technology Based Higher Education and Training (ITHET)*, pages 1–5, Olhao, April 2018. IEEE.
- [23] Mariela Mizota Tamada, José Francisco de Magalhães Netto, and Dhanielly P. R. de Lima. Predicting and reducing dropout in virtual learning using machine learning techniques: A systematic review. *2019 IEEE Frontiers in Education Conference (FIE)*, pages 1–9, 2019.