



Automated Personalized Feedback Improves Learning Gains in An Intelligent Tutoring System

Ekaterina Kochmar^{1,2(✉)}, Dung Do Vu^{1,3}, Robert Belfer¹, Varun Gupta¹,
Iulian Vlad Serban¹, and Joelle Pineau^{1,4}

¹ Korbit Technologies Inc., Montreal, Canada
`ekaterina@korbit.ai`

² University of Cambridge, Cambridge, UK

³ École de Technologie Supérieure, Montreal, Canada

⁴ McGill University & MILA (Quebec Artificial Intelligence Institute),
Montreal, Canada

Abstract. We investigate how automated, data-driven, personalized feedback in a large-scale intelligent tutoring system (ITS) improves student learning outcomes. We propose a machine learning approach to generate personalized feedback, which takes individual needs of students into account. We utilize state-of-the-art machine learning and natural language processing techniques to provide the students with *personalized hints*, *Wikipedia-based explanations*, and *mathematical hints*. Our model is used in *Korbit* (<https://www.korbit.ai>), a large-scale dialogue-based ITS with thousands of students launched in 2019, and we demonstrate that the personalized feedback leads to considerable improvement in student learning outcomes and in the subjective evaluation of the feedback.

Keywords: Intelligent tutoring system · Dialogue-based tutoring system · Natural language processing · Deep learning · Personalized learning and feedback

1 Introduction

Intelligent Tutoring Systems (ITS) [8, 21] attempt to mimic personalized tutoring in a computer-based environment and are a low-cost alternative to human tutors. Over the past two decades, many ITS have been successfully deployed to enhance teaching and improve students' learning experience in a number of domains [1, 2, 5, 6, 9, 12, 17, 19, 22, 23], not only providing feedback and assistance but also addressing individual student characteristics [13] and cognitive processes [27]. Many ITS consider the development of a personalized curriculum and personalized feedback [4, 5, 7, 11, 18, 20, 24, 25], with dialogue-based ITS being some of the most effective tools for learning [3, 14, 15, 21, 26], as they simulate a familiar learning environment of student–tutor interaction, thus helping to improve student motivation. The main bottleneck is the ability of ITS to address




Fig. 1. An example illustrating how the Korbit ITS inner-loop system selects the pedagogical intervention. The student gives an incorrect solution and receives a text hint.

the multitude of possible scenarios in such interactions, and this is where methods of automated, data-driven feedback generation are of critical importance.

Our paper has two major contributions. Firstly, we describe how state-of-the-art machine learning (ML) and natural language processing (NLP) techniques can be used to generate automated, data-driven *personalized hints and explanations*, *Wikipedia-based explanations*, and *mathematical hints*. Feedback generated this way takes the individual needs of students into account, does not require expert intervention or hand-crafted rules, and is easily scalable and transferable across domains. Secondly, we demonstrate that the personalized feedback leads to substantially improved student learning gains and improved subjective feedback evaluation in practice. To support our claims, we utilize our feedback models in Korbit, a large-scale dialogue-based ITS.

2 Korbit Learning Platform

Korbit is a large-scale, open-domain, mixed-interface, dialogue-based ITS, which uses ML, NLP and reinforcement learning to provide interactive, personalized learning online. Currently, the platform has thousands of students enrolled and is capable of teaching topics related to data science, machine learning, and artificial intelligence.

Students enroll based on courses or skills they would like to study. Once a student has enrolled, Korbit tutors them by alternating between short lecture videos and interactive problem-solving. During the problem-solving sessions, the student may attempt to solve an exercise, ask for help, or even skip it. If the student attempts to solve the exercise, their solution attempt is compared against the expectation (i.e. reference solution) using an NLP model. If their solution is classified as incorrect, the inner-loop system (see Fig. 1) will activate and respond with one of a dozen different pedagogical interventions, which include hints, mathematical hints, elaborations, explanations, concept tree diagrams, and multiple choice quiz answers. The pedagogical intervention is chosen by an ensemble of machine learning models from the student's zone of proximal development (ZPD) [10] based on their student profile and last solution attempt.

3 Automatically Generated Personalized Feedback

In this paper, we present experiments on the *Korbit* learning platform with actual students. These experiments involve varying the text hints and explanations based on how they were generated and how they were adapted to each unique student.

Personalized Hints and Explanations are generated using NLP techniques applied by a 3-step algorithm to all expectations (i.e. reference solutions) in our database: (1) keywords, including nouns and noun phrases, are identified within the question (e.g. *overfitting* and *underfitting* in Table 1); (2) appropriate sentence span that does not include keywords is identified in a reference solution using state-of-the-art dependency parsing with *spaCy*¹ (e.g., *A model is underfitting* is filtered out, while *it has a high bias* is considered as a candidate for a hint); and (3) a grammatically correct hint is generated using discourse-based modifications (e.g., *Think about the case*) and the partial hint from step (2) (e.g., *when it has a high bias*).

Table 1. Hint generation. Keywords are marked with boxes

Question	Expectation	Generated hint
What is the difference between overfitting and underfitting ?	A model is underfitting when it has a high bias	<u>Think about the case</u> when it has a high bias

Next, hints are ranked according to their linguistic quality as well as the past student–system interactions. We employ a Random Forest classifier using two broad sets of features: (1) *Linguistic quality features* assess the quality of the hint from the linguistic perspective only (e.g., considering length of the hint/explanation, keyword and topic overlap between the hint/explanation and the question, etc.), and are used by the *baseline model* only. (2) *Performance-based features* additionally take into account past student interaction with the system. Among them, the *shallow personalization model* includes features related to the number of attempted questions, proportion of correct and incorrect answers, etc., and the *deep personalization model* additionally includes linguistic features pertaining to up to 4 previous student–system interaction turns. The three types of feedback models are trained and evaluated on a collection of 450 previously recorded student–system interactions.

Wikipedia-Based Explanations provide alternative ways of helping students to understand and remember concepts. We generate such explanations using another multi-stage pipeline: first, we use a 2 GB dataset on “Machine learning” crawled from Wikipedia and extract all relevant domain keywords from

¹ <https://spacy.io>.

the reference questions and solutions using `spaCy`. Next, we use the first sentence in each article as an *extracted Wikipedia-based explanation* and the rest of the article to *generate candidate explanations*. A Decision Tree classifier is trained on a dataset of positive and negative examples to evaluate the quality of a Wikipedia-based explanation using a number of linguistically-motivated features. This model is then applied to identify the most appropriate Wikipedia-based explanations among the generated ones.

Mathematical Hints are either provided by `Korbit` in the form of suggested equations with gapped mathematical terms for the student to fill in, or in the form of a hint on what the student needs to change if they input an incorrect equation. Math equations are particularly challenging because equivalent expressions can have different representations: for example, y in $y(x + 5)$ could be a function or a term multiplied by $x + 5$. To evaluate student equations, we first convert their \LaTeX string into multiple parse trees, where each tree represents a possible interpretation, and then use a classifier to select the most likely parse tree and compare it to the expectation. Our generated feedback is fully automated, which differentiates `Korbit` from other math-oriented ITS, where feedback is generated by hand-crafted test cases [9, 16].

4 Experimental Results and Analysis

Our preliminary experiments with the *baseline*, *shallow* and *deep personalization* models run on the historical data using 50-fold cross-validation strongly suggested that *deep personalization* model selects the most appropriate personalized feedback. To support our claims, we ran experiments involving 796 annotated student–system interactions, collected from 183 students enrolled for free and studying the machine learning course on the `Korbit` platform between January and February, 2020. First, a hint or explanation was selected at uniform random from one of the personalized feedback models when a student gives an incorrect solution. Afterwards, the student learning gain was measured as the proportion of instances where a student provided a correct solution after receiving a personalized hint or explanation. Since it’s possible for the ITS to provide several pedagogical interventions for a given exercise, we separate the learning gains observed for all students from those for students who received a personalized hint or explanation before their second attempt at the exercise. Table 2 presents the results, showing that the *deep personalization model* leads to the highest student learning gains at 48.53% followed by the *shallow personalization model* at 46.51% and the *baseline model* at 39.47% for all attempts. The difference between the learning gains of the *deep personalization model* and *baseline model* for the students before their second attempt is statistically significant at 95% confidence level based on a z-test ($p = 0.03005$). These results support the hypothesis that automatically generated personalized hints and explanations lead to substantial improvements in student learning gains.

Table 2. Student learning gains for personalized hints and explanations with 95% confidence intervals (C.I.).

Model	All attempts		Before second attempt	
	Mean	95% C.I	Mean	95% C.I
Baseline (No personalization)	39.47%	[24.04%, 56.61%]	37.93%	[20.69%, 57.74%]
Shallow personalization	46.51%	[31.18%, 62.34%]	51.43%	[33.99%, 68.62%]
Deep personalization	48.53%	[36.22%, 60.97%]	60.47%	[44.41%, 75.02%]

Experiments on the *Korbit* platform confirm that extracted and generated *Wikipedia-based explanations* lead to comparable student learning gains. Students rated either or both types of explanations as helpful 83.33% of the time. This shows that automatically-generated Wikipedia-based explanations can be included in the set of interventions used to personalize the feedback. Moreover, two domain experts independently analyzed a set of 86 student–system interactions with *Korbit*, where the student’s solution attempt contained an incorrect mathematical equation. The results showed that over 90% of the *mathematical hints* would be considered either “very useful” or “somewhat useful”.

In conclusion, our experiments strongly support the hypothesis that the personalized hints and explanations, as well as Wikipedia-based explanations, help to improve student learning outcomes significantly. Preliminary results also indicate that the mathematical hints are useful. Future work should investigate how and what types of Wikipedia-based explanations and mathematical hints may improve student learning outcomes, as well as their interplay with student learning profiles and knowledge gaps.

References

1. AbuEl-Reesh, J.Y., Abu-Naser, S.S.: An intelligent tutoring system for learning classical cryptography algorithms (CCAITS). *Int. J. Acad. Appl. Res. (IJAAR)* **2**(2), 1–11 (2018)
2. Agha, M., Jarghon, A., Abu-Naser, S.: An intelligent tutoring system for teaching SQL. *Int. J. Acad. Inf. Syst. Res. (IJAISR)* **2**(2), 1–7 (2018)
3. Ahn, J.W., et al.: Adaptive Visual Dialog for Intelligent Tutoring Systems. In: Rose, C.P., et al. (eds.) *International Conference on Artificial Intelligence in Education*, pp. 413–418. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93846-2_77
4. Al-Dahdooh, R., Abu-Naser, S.: Development and evaluation of the oracle intelligent tutoring system (OITS). *Euro. Acad. Res.* **4**, 8711–8721 (2017)
5. Al-Nakhal, M., Abu-Naser, S.: Adaptive intelligent tutoring system for learning computer theory. *Euro. Acad. Res.* **4**, 8770–8782 (2017)
6. Al Rekhawi, H., Abu-Naser, S.: Android applications UI development intelligent tutoring system. *Int. J. Eng. Inf. Syst. (IJEAIS)* **2**(1), 1–14 (2018)

7. Albacete, P., Jordan, P., Katz, S., Chounta, I.A., McLaren, B.M.: The impact of student model updates on contingent scaffolding in a natural-language tutoring system. In: Isotani, S., Millan, E., Ogan, A., Hastings, P., McLaren, B., Luckin, R. (eds.) International Conference on Artificial Intelligence in Education, pp. 37–47. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-23204-7_4
8. Anderson, J.R., Boyle, C.F., Reiser, B.J.: Intelligent tutoring systems. *Science* **228**(4698), 456–462 (1985)
9. Büdenbender, J., Frischauf, A., Gogvadze, G., Melis, E., Libbrecht, P., Ullrich, C.: Using computer algebra systems as cognitive tools. In: Cerri, S.A., Gouardères, G., Paraguaçu, F. (eds.) ITS 2002. LNCS, vol. 2363, pp. 802–810. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-47987-2_80
10. Cazden, C.: Peekaboo as an Instructional Model: Discourse Development at Home and at School. *Papers and Reports on Child Language Development*, No. 17 (1979)
11. Chi, M., Koedinger, K., Gordon, G., Jordan, P., VanLehn, K.: Instructional factors analysis: a cognitive model for multiple instructional interventions. In: EDM 2011 - Proceedings of the 4th International Conference on Educational Data Mining, pp. 61–70 (2011)
12. Gogvadze, G., Palomo, A.G., Melis, E.: Interactivity of exercises in ActiveMath. In: ICCE, pp. 109–115 (2005)
13. Graesser, A.C., Cai, Z., Morgan, B., Wang, L.: Assessment with computer agents that engage in conversational dialogues and triologues with learners. *Comput. Human Behav.* **76**, 607 – 616 (2017). <https://doi.org/10.1016/j.chb.2017.03.041>, <http://www.sciencedirect.com/science/article/pii/S074756321730198X>
14. Graesser, A.C., Chipman, P., Haynes, B.C., Olney, A.: AutoTutor: an intelligent tutoring system with mixed-initiative dialogue. *IEEE Trans. Educ.* **48**(4), 612–618 (2005)
15. Graesser, A.C., VanLehn, K., Rosé, C.P., Jordan, P.W., Harter, D.: Intelligent tutoring systems with conversational dialogue. *AI Magazine* **22**(4), 39–39 (2001)
16. Hennecke, M.: Online Diagnose in intelligenten mathematischen Lehr-Lern-Systemen. VDI-Verlag (1999)
17. Leelawong, K., Biswas, G.: Designing learning by teaching agents: the Betty’s Brain system. *Int. J. Artif. Intell. Educ.* **18**(3), 181–208 (2008)
18. Lin, C.F., Chu Yeh, Y., Hung, Y.H., Chang, R.I.: Data mining for providing a personalized learning path in creativity: an application of decision trees. *Comput. Educ.* **68**, 199–210 (2013). <https://doi.org/10.1016/j.compedu.2013.05.009>, <http://www.sciencedirect.com/science/article/pii/S0360131513001309>
19. Melis, E., Siekmann, J.: ACTIVE MATH: an intelligent tutoring system for mathematics. In: Rutkowski, L., Siekmann, J.H., Tadeusiewicz, R., Zadeh, L.A. (eds.) ICAISC 2004. LNCS (LNAI), vol. 3070, pp. 91–101. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24844-6_12
20. Munshi, A., Biswas, G.: Personalization in OELEs: developing a data-driven framework to model and scaffold SRL processes. In: Isotani, S., Millan, E., Ogan, A., Hastings, P., McLaren, B., Luckin, R. (eds.) International Conference on Artificial Intelligence in Education. pp. 354–358. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-23207-8_65
21. Nye, B.D., Graesser, A.C., Hu, X.: AutoTutor and family: a review of 17 years of natural language tutoring. *Int. J. Artif. Intell. Educ.* **24**(4), 427–469 (2014)
22. Passier, H., Jeuring, J.: Feedback in an interactive equation solver (2006)
23. Qwaidar, S.R., Abu-Naser, S.S.: Excel intelligent tutoring system. *Int. J. Acad. Inf. Syst. Res. (IJASIR)* **2**(2), 8–18 (2018)

24. Rus, V., Stefanescu, D., Baggett, W., Niraula, N., Franceschetti, D., Graesser, A.C.: Macro-adaptation in conversational intelligent tutoring matters. In: Trausan-Matu, S., Boyer, K.E., Crosby, M., Panourgia, K. (eds.) International Conference on Intelligent Tutoring Systems, pp. 242–247. Springer, Cham (2014) https://doi.org/10.1007/978-3-319-07221-0_29
25. Rus, V., Stefanescu, D., Niraula, N., Graesser, A.C.: DeepTutor: towards macro- and micro-adaptive conversational intelligent tutoring at scale. In: Proceedings of the First ACM Conference on Learning@ Scale Conference, pp. 209–210 (2014)
26. Ventura, M., et al.: Preliminary evaluations of a dialogue-based digital tutor. In: Rose, C. et al. (eds.) International Conference on Artificial Intelligence in Education, pp. 480–483. Springer, Cham (2018) https://doi.org/10.1007/978-3-319-93846-2_90
27. Wu, L., Looi, C.-K.: Agent prompts: scaffolding students for productive reflection in an intelligent learning environment. In: Alevan, V., Kay, J., Mostow, J. (eds.) ITS 2010. LNCS, vol. 6095, pp. 426–428. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13437-1_92