

Question Personalization in an Intelligent Tutoring System

Sabina Elkins^{1,2}, Robert Belfer², Ekaterina Kochmar^{2,3}, Iulian Serban², and Jackie C.K. Cheung^{1,4}

¹ McGill University & MILA (Quebec Artificial Intelligence Institute)

² Korbit Technologies Inc.

³ University of Bath

⁴ Canada CIFAR AI Chair

Abstract. This paper investigates personalization in the field of intelligent tutoring systems (ITS). We hypothesize that personalization in the way questions are asked improves student learning outcomes. Previous work on dialogue-based ITS personalization has yet to address question phrasing. We show that generating versions of the questions suitable for students at different levels of subject proficiency improves student learning gains, using variants written by a domain expert and an experimental A/B test. This insight demonstrates that the linguistic realization of questions in an ITS affects the learning outcomes for students.

Keywords: Intelligent Tutoring System · Dialogue-Based Tutoring System · Personalized Learning

1 Introduction

Intelligent tutoring systems (ITS) are AI systems capable of automating teaching. They have the potential to provide accessible and highly scalable education to students around the world [6]. Previous studies suggest that students learn significantly better in one-on-one tutoring settings than in classroom settings [2]. Personalization can be addressed in an AI-driven, dialogue-based ITSs, and can have significant impact on the learning process [5]. This has been explored in different ways, including dialogue feedback and question selection [8]. To the best of our knowledge, personalization in question phrasing has not been explored.

Students benefit from being asked questions tailored to their level of subject expertise and their needs during in-person tutoring sessions [1,4]. We hypothesize that the same effect can be achieved when questions are adapted to the students' levels of expertise and their needs in an ITS. To test this, we integrate question variants created by a human domain expert onto the Korbit Technologies Inc. platform and run an A/B test. Korbi's AI tutor, *Korbi*, is a dialogue-based ITS, which teaches students by providing them with video lectures and interactive problem solving exercises, selected for each student using ML and NLP techniques [7]. The main contribution of this paper is the demonstration that question personalization in an ITS leads to improvements in learning gains.

2 Methodology

Students interact with *Korbi* through short answer questions and written responses. To assess if the phrasing of these questions can impact learning gains, we first create a set of questions and variants that reformulate the original idea. The variants were created by a human expert from existing questions on the *Korbi* platform. They were designed to reflect three levels of difficulty: *beginner*, *intermediate*, and *advanced*, as per common practice in education. We assume that less detailed questions are harder (as the student must have more background knowledge to understand and answer) and more elaborate ones are easier (as they ‘hint’ at the answer with extra information) [9]. In our data, each question has three variants at different levels of proficiency. Questions were made easier by adding elaborations and synonym replacement, and more difficult by removing non-essential explanations and synonym replacement. As Table 1 shows, the beginner variants are longer and the advanced ones are more concise.




Fig. 1. An question being adapted to different levels while retaining the same answer.

The variants were given to three human experts (data scientists with at least an MSc in a related field) who rated them on three scales from 0 to 5, representing *difficulty* (i.e., the relative complexity of the question as compared to the others), *fluency* (i.e., spelling and grammar), and *meaning preservation* (i.e., if the meaning of the original question is preserved). The mean results of their ratings can be seen in Table 1. The *fluency* and *meaning preservation* metrics are consistently high and the *difficulty* metric increases with the assigned levels.

Table 1. Mean variant scores from human experts, and average word counts by level.

Level	Difficulty	Fluency	Meaning Preservation	Mean Word Count
Beginner	1.689 \pm 0.635	4.600 \pm 0.471	4.789 \pm 0.451	39.800
Intermediate	2.667 \pm 0.689	4.683 \pm 0.481	4.839 \pm 0.406	33.533
Advanced	3.939 \pm 1.269	4.544 \pm 0.661	4.717 \pm 0.516	27.433

The next task is to select an appropriate question variant for each student at each step in the dialogue. Through *Korbi*, we have anonymized access to student history. We isolated 2,137 students’ interactions with the platform. Each student’s history consists of the exercises they encountered and their attempts to solve them. Each exercise encountered was included as a point in our dataset, for a total of 13,504 exercises given to students. Using this, it is possible to calculate a set of features indicative of a student’s level, and subsequently build a logistic regression model to predict if a student will succeed on the next exercise.

The original feature set consisted of 7 features, including overall success rate, improvement (i.e., changes in success rate), skip rates, and others. From this set, two features were selected based on their contribution to the best model in the preliminary experiments: (1) *topic success feature* is a numerical feature in $[0, 1]$ that shows the eventual success rate per all exercises previously attempted in a given topic, and (2) *topic skip feature* is a numerical feature in $[0, 1]$ that is the skip rate per all exercises previously attempted in a given topic. A topic on Korbi is a broad category of material, such as ‘Probability’. Using these features the model is able to predict next exercise success with an accuracy of 80%.

The variant assignment model calculates the features when a student gets a new exercise, and generates a probability of success with the regression model. Students are assigned variants based on the percentile range that their probability of success falls into (calculated from the predictions across the entire data set). Students in the 0th to 33rd percentiles get beginner variants, in the 33rd to 66th percentiles get intermediate variants, and the rest get advanced ones.

3 Results and Analysis

To test our claims, we put the variants and assignment model described in Section 2 on the Korbi platform. Our A/B test ran over 2 months, collecting data from over 400 students at varied skill levels. Student attempts were divided into three groups. The *expected* variant group received the variant which matched their assignment model score. The *non-expected* variant group received a variant which did not match their score from the assignment model (e.g., beginner question for an advanced student). The *control* group students received the original variant (i.e., that which was already on the platform before this experiment).

Table 2. Test results. Metrics marked with * are statistically significant at the $\alpha = 0.05$ level by a Student’s *t*-test.

Experiment Group	Solution Acceptance*	Ultimate Failure Rate*	Skip Rate	n
Expected	0.626 ± 0.069	0.163 ± 0.053	0.105 ± 0.044	190
Non-Expected	0.468 ± 0.083	0.295 ± 0.076	0.144 ± 0.058	139
Control	0.596 ± 0.081	0.191 ± 0.065	0.121 ± 0.054	141

Solution acceptance is the proportion of success per exercise attempts. However, succeeding on exercises does not equate to learning. Students should be challenged within their zone of proximal development [3] but eventually obtain the right answer, so we aim to minimize the *ultimate failure rate* as opposed to simply maximizing attempt success. This metric is the proportion of failure out of all exercises seen by students. Unlike *solution acceptance* which shows the success rate per attempt, *ultimate failure rate* shows the fail rate per exercise. *Skip rate* is indicative of a student’s engagement. Intuitively, the more they skip, the less they engage with the content. All three of these metrics show the *expected* group performing the best, followed by *control* and finally *non-expected*. For *solution acceptance* and *ultimate failure rate*, the difference between *expected* and *non-expected* groups is statistically significant at $\alpha = 0.05$ by a Student’s *t*-test.

The difference between the *expected* and *control* groups is smaller than the difference between the *expected* and *non-expected* groups. This can be attributed to the fact that the original questions were refined through several rounds of review by domain experts when they were created for Korbi platform, whereas the variants only were reviewed once. Additionally, the *control* group’s exercises are always intermediate or advanced, while the strongest result is seen with beginners. Isolating the students who score for beginner variants only, we see a 19% relative reduction in *ultimate failure* when comparing the *expected* to the *control* group, which demonstrates a bigger impact for beginners. Additionally, the same comparison shows a 30% relative reduction in the *skip rate*, suggesting that the beginners are more engaged when dealing with beginner variants.

4 Conclusion

We see a clear improvement in the success of students in the *expected* group. This confirms our hypothesis that providing question variants suited to student’s level improves their learning gains. These variants are more useful for beginner students who need more assistance, which is an encouraging and intuitive result. The future of this work is in automating the creation question variants for scalability, and creating a more sophisticated variant assignment approach.

Acknowledgements We’d like to thank Korbit for hosting our experiment on their platform, and Mitacs for their grant to support this project. We are grateful to the anonymous reviewers for their valuable feedback.

References

1. Ashton-Jones, E. (1988). Asking the Right Questions: A Heuristic for Tutors. The Writing Center Journal, 9(1), 29-36.
2. Bausell, R., Moody, W., & Walzl, F. (1972). A factorial study of tutoring versus classroom instruction. American Educational Research Journal, 9(4), 591-597.
3. Cazden, C. (1979). Peekaboo as an Instructional Model: Discourse Development at Home and at School. Papers and Reports on Child Language Development, No. 17.
4. Hrastinski, S., et al. (2019). Identifying and exploring the effects of different types of tutor questions in individual online synchronous tutoring in mathematics. Interactive Learning Environments, 1-13.
5. Kochmar E., et al. (2020). Automated Personalized Feedback Improves Learning Gains in An Intelligent Tutoring System. AIED 2020, LNAI 12164, pp. 140–146.
6. Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of Intelligent Tutoring Systems: A Meta-Analytic Review. Review of Educational Research, 86(1), 42–78.
7. Serban, Iulian Vlad et al. (2020). A large-scale, open-domain, mixed-interface dialogue-based ITS for STEM. AIED 2020, 387-392.
8. St-Hilaire, F. et al. (2022). A New Era: Intelligent Tutoring Systems Will Transform Online Learning for Millions. arXiv:2203.03724
9. Taylor, R.S. (1962). The process of asking questions. Amer. Doc., 13: 391-396.