



# A Large-Scale, Open-Domain, Mixed-Interface Dialogue-Based ITS for STEM

Iulian Vlad Serban<sup>1(✉)</sup>, Varun Gupta<sup>1</sup>, Ekaterina Kochmar<sup>1,2</sup>, Dung D. Vu<sup>1,3</sup>,  
Robert Belfer<sup>1</sup>, Joelle Pineau<sup>1,4</sup>, Aaron Courville<sup>1,4</sup>, Laurent Charlin<sup>1,4</sup>,  
and Yoshua Bengio<sup>1,4</sup>

<sup>1</sup> Korbit Technologies Inc., Montreal, Canada  
iulian@korbit.ai

<sup>2</sup> University of Cambridge, Cambridge, UK

<sup>3</sup> École de Technologie Supérieure, Montreal, Canada

<sup>4</sup> MILA (Quebec Artificial Intelligence Institute), Montreal, Canada

**Abstract.** We present *Korbit*, a large-scale, open-domain, mixed-interface, dialogue-based intelligent tutoring system (ITS). *Korbit* uses machine learning, natural language processing and reinforcement learning to provide interactive, personalized learning online. *Korbit* has been designed to easily scale to thousands of subjects, by automating, standardizing and simplifying the content creation process. Unlike other ITS, a teacher can develop new learning modules for *Korbit* in a matter of hours. To facilitate learning across a wide range of STEM subjects, *Korbit* uses a mixed-interface, which includes videos, interactive dialogue-based exercises, question-answering, conceptual diagrams, mathematical exercises and gamification elements. *Korbit* has been built to scale to millions of students, by utilizing a state-of-the-art cloud-based micro-service architecture. *Korbit* launched its first course in 2019 and has over 7,000 students have enrolled. Although *Korbit* was designed to be open-domain and highly scalable, A/B testing experiments with real-world students demonstrate that both student learning outcomes and student motivation are substantially improved compared to typical online courses.

**Keywords:** Intelligent tutoring system · Dialogue-based tutoring system · Natural language processing · Reinforcement learning · Deep learning · Personalized · Interactive learning · Data science · STEM

## 1 Introduction

Intelligent tutoring systems (ITS) are computer programs powered by artificial intelligence (AI), which deliver real-time, personalized tutoring to students. Traditional ITS implement or imitate the behavior and pedagogy of human tutors. In particular, one type of ITS are dialogue-based tutors, which use natural language conversations to tutor students [13]. This process is sometimes

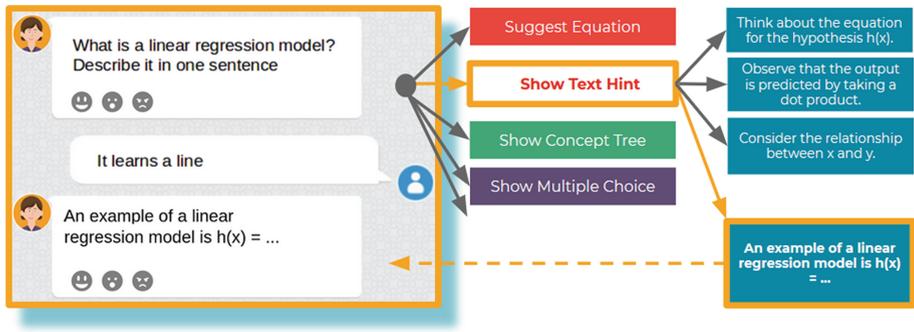
called “Socratic tutoring”, because of its similarity to Socratic dialogue [17]. Newer ITS have started to interleave their dialogue with interactive media (e.g. interactive videos and web applets) – a so-called “mixed-interface system”. It has been shown that ITS can be twice as effective at promoting learning compared to the previous generation of computer-based instruction and may be as effective as human tutors in general [12].

However, despite the fact that ITS have been around for decades and are known to be highly effective, their deployment in education and industry has been extremely limited [14, 16]. A major reason for this is the sheer cost of development [5, 14]. As observed by Olney [14]: *“Unfortunately, ITS are extremely expensive to produce, with some groups estimating that it takes 100 h of authoring time from AI experts, pedagogical experts, and domain experts to produce 1 h of instruction.”* On the other hand, lower-cost educational approaches, such as massive open online courses (MOOCs), have flourished and now boast of having millions of learners. It is estimated that today there are over 110 million learners around the world enrolled in MOOCs [18]. However, the learning outcomes resulting from learning in MOOCs depend critically on their teaching methodology and quality of content, and remains questionable in general [2, 3, 9–11, 15]. In particular, recent research indicates that MOOCs having low levels of active learning, little feedback from instructors and peers, and few peer discussions tend to yield poor learning outcomes [10, 15]. Further, it is well-known that student retention in MOOCs is substantially worse than in traditional classroom learning [8]. By combining low cost and scalability with the personalization and effectiveness of ITS, we hope *Korbit* may help to effectively teach and motivate millions of students around the world.

## 2 The *Korbit* ITS

*Korbit* is a large-scale, open-domain, mixed-interface, dialogue-based ITS, which uses machine learning, natural language processing (NLP) and reinforcement learning (RL) to provide interactive, personalized learning online. The ITS has over 7,000 students enrolled from around the world, including students from educational institutions and professionals from industry partners. *Korbit* is capable of teaching topics related to data science, machine learning, and artificial intelligence. The modular platform will soon be expanded with many more topics.

Students enroll on the *Korbit* website by selecting either a course or a set of skills they would like to study. Students may also answer a few questions about their background knowledge. Based on these, *Korbit* generates a personalized curriculum for each student. Following this, *Korbit* tutors the student by alternating between short lecture videos and interactive problem-solving exercises. The *outer-loop* system decides on which lecture video or exercise to show next based on the personalized curriculum. Work is currently underway to adapt the curriculum during the learning process (Fig. 1).



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

During the exercise sessions, the *inner-loop* system manages the interaction. First, it shows the student a problem statement (e.g., a question). The student may then attempt to solve the exercise, ask for help, or skip the exercise. 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, then the *inner-loop* system will select one of a dozen different pedagogical interventions. The pedagogical interventions include textual 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 based on the student's profile and last solution attempt. Depending on the pedagogical intervention, the *inner-loop* system may either ask the student to retry the initial exercise or follow up on the intervention (e.g., with additional questions, confirmations, or prompts).

The *KorbIt* ITS is related to the work on dialogue-based ITS, such as the pioneering *AutoTutor* and the newer *IBM Watson Tutor* [1, 6, 7, 13, 19]. Although *KorbIt* is highly constrained compared to existing dialogue-based ITS, a major innovation of *KorbIt* lies in its modular, scalable design. The *inner-loop* system is implemented as a finite-state machine. Each pedagogical intervention is a separate state, with its own logic, data and machine learning models. Each state operates independently of the rest of the system, has access to all database content (including exercises and videos) and can autonomously improve as new data becomes available. This ensures that the system gets better and better, that it can adapt to new content and that it can be extended with new pedagogical interventions. The transitions between the states of the finite-state machine is decided by a reinforcement learning model, which itself is agnostic to the underlying implementation of each state and also continues to improve as more and more data becomes available.

### 3 System Evaluation

We have conducted multiple studies to evaluate the *Korbit* ITS. Some of these studies have evaluated the entire system while others have focused on particular aspects or modules of the system. Taken together, the studies demonstrate that the *Korbit* ITS is an effective learning tool and that it overall improves student learning outcomes and motivation compared to alternative online learning approaches.

In this paper we limit ourselves and discuss only one of these studies. The study we present compares the entire system (**Full ITS**) against an xMOOC-like system [4]. The purpose of this particular study is to evaluate 1) whether students prefer the *Korbit* ITS or a regular MOOC, 2) whether the *Korbit* ITS increases student motivation, and 3) which aspects of the *Korbit* ITS students find most useful and least useful. In an ideal world, *Korbit* ITS would be compared against a regular xMOOC teaching students through lecture videos and multiple choice quizzes in a randomized controlled trial (a randomized A/B testing experiment). However, it is not possible to compare against such a system in a randomized controlled trial, because it would create confusion and drastically offset student expectations. Therefore, in this study, we compare the **Full ITS** against a reduced ITS, which appears identical to the **Full ITS** and utilizes the same content (video lectures and exercise questions), but defaults to multiple choice quizzes 50% of the time. Thus, students assigned to the reduced ITS spend about half of their interactions in an xMOOC-like setting. We refer to this system as the **xMOOC ITS**.

**Table 1.** A/B testing results comparing the **Full ITS** against the **xMOOC ITS**: average time spent by students (in minutes), returning students (in %), students who said they will refer others (in %) and learning gain (in %), with corresponding 95% confidence intervals. The \* and \*\* shows statistical significance at 90% and 95% confidence level.

System	Time spent	Returning students	Will refer others	Learning gain
xMOOC ITS	22.98 ± 4.18	26.98% ± 3.44%	44.83% ± 9.00%	<b>39.14% ± 2.35%</b>
Full ITS	<b>39.86 ± 3.70**</b>	<b>31.69% ± 1.92*</b>	<b>54.17% ± 4.05%</b>	

The experiment was conducted in 2019 with  $n = 612$  participants. Students who enrolled online were randomly assigned to either the **Full ITS** (80%) or **xMOOC ITS** (20%). Students came from different countries and were not subject to any selection or filtering process. Apart from bug fixes and speed improvements, the system was not modified during the experiment to limit confounding factors. After studying for about 45 min, students were shown a questionnaire to evaluate the system.

Table 1 shows the experimental results. The average time spent in the **Full ITS** was 39.86 min compared to 22.98 min in the **xMOOC ITS**. As such, the **Full ITS** yields a staggering 73.46% increase in time spent. In addition, the percentage

of returning students and the percentage of students who said they would refer others to use the system is substantially higher for the Full ITS compared to the xMOOC ITS. These results were also confirmed by the feedback provided by the students in the questionnaire. Thus, we can conclude that students strongly prefer Korbit ITS over xMOOCs and that the Korbit ITS increases overall student motivation.

Table 1 also shows that the average student learning was observed to be 39.14%. The learning gain is measured as the proportion of instances where a student provides a correct exercise solution after having receiving a pedagogical intervention from the Korbit ITS. Thus, the pedagogical interventions appear to be effective.

Finally, in the questionnaire, 85.31% of students reported that they found the chat equally or more fun compared to learning alone and 66.67% of students reported that the chat helped them learn better sometimes, many times or all of the time. For the Full ITS, 54.17% of students reported that they would refer others to use Korbit ITS. In addition, students reported that the Korbit ITS could be improved by more accurately identifying their solutions as being correct or incorrect and, in the case of incorrect solutions, by providing more personalized feedback.

## References

1. Ahn, J.-W., et al.: Adaptive visual dialog for intelligent tutoring systems. In: Penstein Rosé, C., et al. (eds.) AIED 2018. LNCS (LNAI), vol. 10948, pp. 413–418. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-93846-2\\_77](https://doi.org/10.1007/978-3-319-93846-2_77)
2. Cavanaugh, J.K., Jacquemin, S.J.: A large sample comparison of grade based student learning outcomes in online vs. face-to-face courses. *Online Learn.* **19**(2), n2 (2015)
3. Colvin, K.F., Champaign, J., Liu, A., Zhou, Q., Fredericks, C., Pritchard, D.E.: Learning in an introductory physics MOOC: all cohorts learn equally including an on-campus class. *Int. Rev. Res. Open Distrib. Learn.* **15**(4), 263–283 (2014)
4. Daniel, J.: Making sense of MOOCs: Musings in a maze of myth, paradox and possibility. *J. Interact. Media Educ.* **2012**(3), 18 (2012)
5. Folsom-Kovarik, J.T., Schatz, S., Nicholson, D.: Plan ahead: pricing ITS learner models. In: Proceedings of the 19th Behavior Representation in Modeling & Simulation (BRIMS) Conference, pp. 47–54 (2010)
6. 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)
7. Graesser, A.C., VanLehn, K., Rosé, C.P., Jordan, P.W., Harter, D.: Intelligent tutoring systems with conversational dialogue. *AI Mag.* **22**(4), 39–39 (2001)
8. Hone, K.S., El Said, G.R.: Exploring the factors affecting mooc retention: A survey study. *Comput. Educ.* **98**, 157–168 (2016)
9. Kirtman, L.: Online versus in-class courses: an examination of differences in learning outcomes. *Issues Teach. Educ.* **18**(2), 103–116 (2009)

10. Koedinger, K.R., Kim, J., Jia, J.Z., McLaughlin, E.A., Bier, N.L.: Learning is not a spectator sport: Doing is better than watching for learning from a MOOC. In: Proceedings of the Second (2015) ACM Conference on Learning@ Scale, pp. 111–120 (2015)
11. Koxvold, I.: MOOCs: opportunities for their use in compulsory-age education. Department for Education (2014)
12. Kulik, J.A., Fletcher, J.: Effectiveness of intelligent tutoring systems: a meta-analytic review. *Rev. Educ. Res.* **86**(1), 42–78 (2016)
13. 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)
14. Olney, A.M.: Using novices to scale up intelligent tutoring systems. In: Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) (2018)
15. Otto, D., Bollmann, A., Becker, S., Sander, K.: It's the learning, stupid! discussing the role of learning outcomes in moocs. *Open Learn. J. Open Dist. e-Learn.* **33**(3), 203–220 (2018)
16. Ritter, S., Anderson, J.R., Koedinger, K.R., Corbett, A.: Cognitive tutor: applied research in mathematics education. *Psychon. Bull. Rev.* **14**(2), 249–255 (2007)
17. Rosé, C.P., Moore, J.D., VanLehn, K., Allbritton, D.: A comparative evaluation of socratic versus didactic tutoring. In: Proceedings of the Annual Meeting of the Cognitive Science Society, vol. 23 (2001)
18. Shah, D.: By The Numbers: MOOCs in 2019. Class Central MOOC Report (2019)
19. Ventura, M., et al.: Preliminary evaluations of a dialogue-based digital tutor. In: Penstein Rosé, Penstein Rosé, C., et al. (eds.) AIED 2018. LNCS (LNAI), vol. 10948, pp. 480–483. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-93846-2\\_90](https://doi.org/10.1007/978-3-319-93846-2_90)