THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Beetle-Grow: An Effective Intelligent Tutoring System for Data Collection

Citation for published version:

Farrow, E, Dzikovska, MO & Moore, JD 2016, Beetle-Grow: An Effective Intelligent Tutoring System for Data Collection. in *L @S '16 Proceedings of the Third (2016) ACM Conference on Learning @ Scale*. ACM, pp. 177-180, Third (2016) ACM Conference on Learning @ Scale, Edinburgh, United Kingdom, 25/04/16. <https://doi.org/10.1145/2876034.2893408>

Digital Object Identifier (DOI):

[10.1145/2876034.2893408](https://doi.org/10.1145/2876034.2893408)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

L@S '16 Proceedings of the Third (2016) ACM Conference on Learning @ Scale

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Beetle-Grow: An Effective Intelligent Tutoring System for Data Collection

Elaine Farrow

University of Edinburgh
School of Informatics
10 Crichton Street
Edinburgh, EH8 9AB, UK
elaine.farrow@ed.ac.uk

Myroslava O. Dzikovska¹

University of Edinburgh
School of Informatics
10 Crichton Street
Edinburgh, EH8 9AB, UK
m.dzikovska@ed.ac.uk

Johanna D. Moore

University of Edinburgh
School of Informatics
10 Crichton Street
Edinburgh, EH8 9AB, UK
j.moore@ed.ac.uk

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the owner/author(s).
L@S 2016, April 25–26, 2016, Edinburgh, Scotland, UK.
ACM 978-1-4503-3726-7/16/04.
<http://dx.doi.org/10.1145/2876034.2893408>

Abstract

We present the BEETLE-GROW intelligent tutoring system, which combines active experimentation, self-explanation, and formative feedback using natural language interaction. It runs in a standard web browser and has a fresh, engaging design. The underlying back-end system has previously been shown to be highly effective in teaching basic electricity and electronics concepts.

BEETLE-GROW has been designed to capture student interaction and indicators of learning in a form suitable for data mining, and to support future work on building tools for interactive tutoring that improve after experiencing interaction with students, as human tutors do.

Author Keywords

intelligent tutoring; natural language; interaction data; conceptual learning; physics; electronics

Introduction

Many effective teaching techniques combine active experimentation, self-explanation, and formative feedback [1, 2, 7, 8, 9, 10]. In classroom situations, these rely heavily on dialogue with a teacher. Our work focuses on implementing similar techniques in an Intelligent Tutoring System (ITS) using natural language interaction.

¹Myroslava Dzikovska is now at Google, London

Figure 1: The graphics used in the interactive circuit workbench combine selective stylization with an element of realism to facilitate a constructivist learning approach [3], encouraging learners to experiment freely and learn from hands-on experience.

Figure 2: The circuit workbench communicates with a back-end circuit simulator to show the correct voltmeter readings and display each bulb at the appropriate level of brightness.

In learning environments with a very high student-teacher ratio such as MOOCs, it is clearly not practical for every learner to engage in extended dialogue with an expert human tutor. A dialogue-based ITS can promote learning by guiding experimentation, eliciting predictions and explanations, and generating tailored formative feedback automatically for a large population of learners.

Automated methods for analyzing student input often use statistical models trained on large corpora of student language [6]. We would like to go beyond identifying correct and incorrect student answers in text and start to use interaction data to learn effective feedback tactics. This would allow a future ITS to improve from interaction with students, just as human tutors do [11]. A necessary first step is to collect the required training data.

The BEETLE-GROW system we describe in this paper is designed for large-scale data collection, while also delivering a valuable learning experience in its own right. An interactive demonstration of the system can be accessed at <http://beetle-grow.inf.ed.ac.uk>.

Background

BEETLE-GROW builds on the success of BEETLE II, a natural-language ITS which has been shown to be highly effective in teaching basic electricity and electronics, demonstrating effect sizes of $d=1.72$ [5]. BEETLE II combines natural language dialogue, guided experimentation, and adaptive feedback to implement the conceptual change [4] tutoring approach.

In BEETLE-GROW, we have added a new browser-based user interface, allowing multiple remote users to access the system concurrently. Since the only technical requirement is an HTML5 compliant browser (such as Google Chrome or Mozilla Firefox), the system can easily be made

available to a wide variety of students. Student interaction data is logged in a format suitable for data mining and as training data for machine learning.

Figure 3: The browser-based user interface uses industry standard technologies: HTML5, CSS3, JavaScript, and the Play Framework, an open source web programming framework.

User Interface

The user interface (Figure 3) uses a responsive design optimized for modern browsers, including touch screen devices such as Android and iPad tablets. It is divided into panels corresponding to different aspects of the tutoring process: reading learning materials (Figure 4), interacting with the tutor (Figure 5), and experimenting with electrical circuits using a circuit simulator (Figures 1 and 2). The layout is similar to the BEETLE II interface, but redesigned with a friendly and contemporary look and feel, intended to be attractive to our target user group (students aged around 13-25).

Figure 5: Student input is visually separated from tutor feedback and questions in the chat interface, making the dialogue easier to follow.

Figure 4: The learning materials are presented using HTML5, with scalable text and high quality graphics optimized for presentation in modern web browsers.

System Architecture

BEETLE-GROW is implemented as a distributed system, in which a web server accepts requests from the GUI and routes users to the correct page. Lesson progress is tracked in a database, allowing the student to exit the lesson at any time and resume later from the start of the last exercise attempted.

The curriculum covers ten main topics and includes over 150 questions. Some require the student to select and connect electrical components in the interactive circuit workbench to match a circuit diagram shown on screen. Others expect a written answer of one or two sentences. There are numerical and yes/ no questions, and questions asking the student to identify particular circuit components.

Student input consists of circuits built in the workbench, text entered in the chat window, and button presses to request help or to move through the learning materials. The dialogue manager is responsible for processing the

input and deciding on the next action to take, such as presenting further learning materials, asking a question, or giving feedback on a student answer. It communicates with the back-end components retained from BEETLE II, which handle curriculum management, natural language processing, answer diagnosis, and tutorial planning.

The tutorial planner does not allow a student to move on from a question without supplying an answer in the expected form: either a circuit or some text. Submitted circuits are compared against the desired circuit answer, a diagnosis is made, and formative feedback is given; for example, "There's a problem with the circuit. The switch should not be open." After a third successive failure, the correct circuit is loaded into the workbench, allowing the student to proceed.

In contrast, no formative feedback is currently given for textual answers. Instead, the system simply states the expected answer and leaves the student to decide if their own answer was a close enough match. This behavior corresponds to the TELL condition reported in [5], which demonstrated that asking students for an explanation and then immediately giving away the correct answer was highly effective, despite the obvious potential to 'game the system' by typing minimal or nonsense answers. This approach also avoids the negative effect on student satisfaction and engagement seen in BEETLE II evaluations due to the limitations of the natural language understanding and feedback selection modules used there.

Conclusion and Future Work

BEETLE-GROW harnesses the proven educational potential of the BEETLE II system and makes it available to a new audience. The browser-based user interface makes deployment at scale straightforward. We have

confidence that BEETLE-GROW will help students to learn, while simultaneously allowing us to collect a large corpus of interaction data for future analysis.

As our next step, we are interested in partnering with teachers and other education researchers to carry out large-scale user trials with BEETLE-GROW in the classroom and remotely. Data from such trials will generate new insights into the effect of active experimentation on learning in the domain of basic electricity and electronics.

Student answers and interaction data collected from these trials can also be used to train statistical algorithms, such as those described in [6], in order to provide better student feedback in a future version of BEETLE-GROW.

Acknowledgments

We thank the US Office of Naval Research for their financial support under Award number N000141410733, and our colleagues Mary Hutchison and Edmund Farrow for their contributions to the design and implementation of the new user interface.

References

- [1] Chi, M., Jordan, P., VanLehn, K., and Litman, D. To elicit or to tell: Does it matter? In *Proceedings 14th International Conference on Artificial Intelligence in Education (AIED)* (Brighton, UK, 2009).
- [2] Chi, M. T. H., de Leeuw, N., Chiu, M.-H., and LaVancher, C. Eliciting self-explanations improves understanding. *Cognitive Science* 18, 3 (1994), 439–477.
- [3] Duffy, T. M., and Jonassen, D. H. *Constructivism and the technology of instruction: A conversation*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1992.
- [4] Duit, R., and Treagust, D. F. Conceptual change: A powerful framework for improving science teaching and learning. *International Journal of Science Education* 25, 6 (2003), 671–688.
- [5] Dzikovska, M., Steinhauser, N., Farrow, E., Moore, J., and Campbell, G. Beetle II: Deep natural language understanding and automatic feedback generation for intelligent tutoring in basic electricity and electronics. *International Journal of Artificial Intelligence in Education* 24, 3 (9 2014), 284–332.
- [6] Dzikovska, M. O., Nielsen, R., Brew, C., Leacock, C., Giampiccolo, D., Bentivogli, L., Clark, P., Dagan, I., and Dang, H. T. SemEval-2013 Task 7: The Joint Student Response Analysis and 8th Recognizing Textual Entailment Challenge. In *Proceedings of the 6th International Workshop on Semantic Evaluation (SEMEVAL-2013)*, Association for Computational Linguistics (Atlanta, Georgia, USA, 2013).
- [7] Litman, D., Moore, J., Dzikovska, M., and Farrow, E. Using natural language processing to analyze tutorial dialogue corpora across domains and modalities. In *Proceedings of 14th International Conference on Artificial Intelligence in Education (AIED)* (Brighton, UK, 2009).
- [8] Narciss, S. The impact of informative tutoring feedback and self-efficacy on motivation and achievement in concept learning. *Experimental Psychology* 51, 3 (2004), 214–228.
- [9] Purandare, A., and Litman, D. Content-learning correlations in spoken tutoring dialogs at word, turn and discourse levels. In *Proceedings of the 21st International FLAIRS Conference* (2008).
- [10] Shute, V. J. Focus on formative feedback. *Review of educational research* 78, 1 (2008), 153–189.
- [11] VanLehn, K. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist* 46, 4 (10 2011), 197–221.