Using Data to Explore the Differences between Instructional Vision and Student Performance

Erik Harpstead
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15232 USA
eharpste@cs.cmu.edu

Vincent Aleven
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15232 USA
aleven@cs.cmu.edu

Christopher J. MacLellan
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15232 USA
cmaclell@cs.cmu.edu

Kenneth R. Koedinger
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15232 USA
koedinger@cs.cmu.edu

Abstract
In any instructional setting it is important to take into account the ways in which a student's experience differs from the one that was envisioned by the instructional designer. When providing instruction in a massive online context, the differences between an instructor's vision and a student's experience is amplified. We have been exploring techniques that allow designers to tease apart the idiosyncrasy of their instructional interventions in order to improve them. In this presentation we will cover two techniques that we have employed to evaluate and better understand instructional interventions.

Author Keywords
Educational Technology, Learning Curves, Clustering Student Solutions

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous; K.3.1 Computer Uses in Education: Computer-assisted instruction (CAI).

Introduction
When designing technological instruction it is important to keep in mind that the experience of an individual

student may not correspond with the instructional designer's intention. This issue can become compounded when the intervention is taken to scale. In such cases it is important to have tools and techniques capable of helping instructional designers understand how student performance differs from their expectation and, ideally, provide them with a direction as to how to mitigate the difference. We and our other colleagues in the LearnLab have explored some techniques for looking at this issue that we believe could benefit the growing community of massive online education. The two techniques we wish to highlight are Learning Curve Analysis and Solution Space Clustering.

Learning Curve Analysis

Understanding an instructional intervention requires a robust and detailed description of student behavior. To make analysis easier all of the educational technologies developed as part of the LearnLab are designed to work with the PSLC DataShop [5]. The PSLC DataShop is the largest repository of educational technology interaction data in the world, storing over 111,000 hours of student data from more than 50 projects. In addition to storing student data the DataShop is also equipped with a number of tools to facilitate analysis of instructional interventions.

One of DataShop's most powerful capacities is its ability to calculate learning curves of student performance on target skills, such as the ones in Figure 1. These curve show the error rate of student performance compared to the number of opportunities they have had to demonstrate mastery. The theory behind the use of learning curves is called the "power law of learning," which suggests that error rates on a given skill decrease following a power law as the skill is learned [1]. These learning curves can be used to compare multiple instructional interventions by showing which one leads to mastery in fewer opportunities.

In addition to serving as a measurement of learning, learning curves can be used to expose latent skill demands that might not have been apparent to the

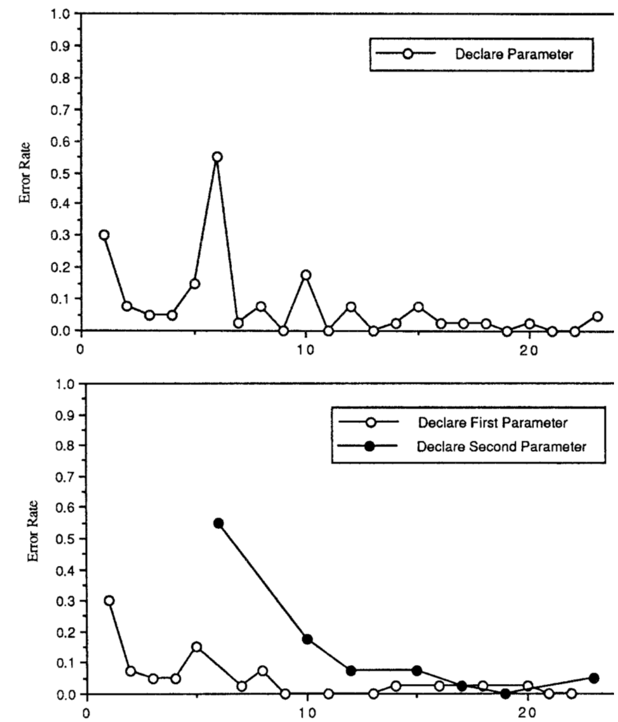

Figure 1. Two examples of learning curves of student performance on skills in a programing tutor showing the decrease of error rate over opportunities. The top curve shows an anomalous blip caused by the conflation of two separate skills as one. The lower curve demonstrates the effect of splitting the single skill into two separate ones.

designer. The top curve in Figure 1 is taken from a programming tutor where the designers coded a skill for declaring a parameter of a function [2]. The curve deviates from the expected shape with a large blip in the middle. This blip was caused by the fact that a number of the problems required students to declare more than 1 parameter and the addition of extra syntactic troubles lead them to make errors again. When the model is changed to account for declaring single and multiple parameters as different skills then the power law shape becomes apparent for each skill individually. Discovering patterns like this can help designers to understand if they need to introduce new instruction to address overlapping skills that were previously being considered the same. DataShop provides tools to assist designers in discovering these types of anomalies automatically in their own data [6].

Solution Space Clustering

While learning curve analysis is a powerful technique it does require a mapping of student actions to particular skills. Creating such a mapping requires a broad understanding of what students are capable of doing in your system, which can be difficult on more open-ended tasks such as designing a product that satisfies a set of constraints. In such domains it is still important to understand how student performance might differ from designer expectation. To address this issue we have done work in mapping out the space of student solutions in an open-ended educational game using feature extraction and clustering.

Our approach to cluster student solutions employs a technique called Conceptual Feature Extraction [4]. At a high level, the technique works by viewing student solutions as a parsing problem. It first induces a grammar to describe the solutions and the uses the grammar rules to parse solutions and generate feature vectors suitable for clustering. Each cluster is then interpreted as one of the possible solutions to that level.

Using the clusters as a mapping of possible solutions to the level, designers can then see which solutions match the one that they had in mind and look at the relative frequency of students creating that solution compared to any other one. This data generates a graph similar to Figure 2. This plot shows the percentage of students who created the envisioned solution to each level as opposed to any other solution. While the plot is sorted by percentage the levels to the left do tend to be earlier tutorial levels while levels to the right tend to be more complicated later levels of the game. Armed with this information designers can then inspect what is going on in the levels with small percentages of usage and from there decide what intervention is necessary, e.g.

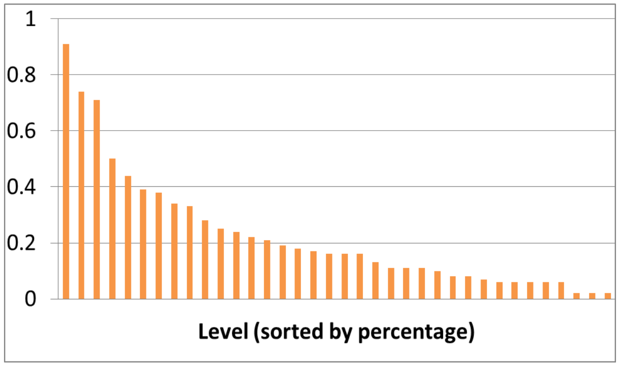

Figure 2. A graph of the percentage of students who used the designer envisioned solution to each level of an educational game. The levels are sorted by percentage with mostly envisioned solutions to the left.

changing the victory condition of the game to be more or less permissive of a particular type of design. We have recently been exploring ways of further informing this process using heuristics of goal and feedback alignment [3].

Conclusion

As instruction moves increasingly toward a massive online context it will become more important for instructional designers to understand where their interventions might break down. We hope our techniques might inspire others to be wary of these issues and find ways to mitigate them themselves.

References

[1] Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., and Qin, Y. An integrated theory of the mind. *Psychological review 111*, 4 (2004), 1036–60.

[2] Corbett, A.T. and Anderson, J.R. Knowledge Tracing : Modeling the Acquisition of Procedural Knowledge. *User Modelling and User-Adapted Interaction 4*, (1995), 253–278.

[3] Harpstead, E., Maclellan, C.J., Aleven, V., and Myers, B.A. Using Extracted Features to Inform Alignment-Driven Design Ideas in an Educational Game. *Proc. CHI '14*, (2014), Submitted.

[4] Harpstead, E., Maclellan, C.J., Koedinger, K.R., Aleven, V., Dow, S.P., and Myers, B.A. Investigating the Solution Space of an Open-Ended Educational Game Using Conceptual Feature Extraction. *Proc. EDM 2013*, (2013).

[5] Koedinger, K.R., Baker, R.S.J. d, Cunningham, K., Skogsholm, A., Leber, B., and Stamper, J. A Data Repository for the EDM community: The PSLC DataShop. In C. Romero, S. Ventura, M. Pechenizkiy and R.S.J. d. Baker, eds., *Handbook of Educational Data Mining*. 2010, 43–55.

[6] Koedinger, K.R., Mclaughlin, E.A., and Stamper, J.C. Automated Student Model Improvement. *Proc. EDM 2012*, (2012).