Learning the Features Used To Decide How to Teach

Min Hyung Lee
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh PA 15213 USA
minhyunl@andrew.cmu.edu

Zhuoying Wang
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh PA 15213 USA
zhuoyinw@andrew.cmu.edu

Joe Runde
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh PA 15213 USA
jrunde@andrew.cmu.edu

Emma Brunskill
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh PA 15213 USA
ebrun@cs.cmu.edu

Warfa Jibril
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh PA 15213 USA
wjibril@andrew.cmu.edu

L@S 2015, Mar 14-18, 2015, Vancouver, BC, Canada
ACM 978-1-4503-3411-2/15/03.
http://dx.doi.org/10.1145/2724660.2728707

Abstract
As a step towards scaling personalized instruction, we seek to automatically identify the key features of the interactive learning process teachers use to select the next activity when teaching a single student. Such features could both inform computational student models designed to facilitate instructional decisions, and help enable automated self-improving teaching systems that leverage this identified feature set. We present preliminary results that a very small set of features is almost as good as a much larger set of features at predicting human tutor decisions when teaching students about histograms.

Introduction
Massive open online classes (MOOCs) are transforming access to education, enabling 1 or a few teaching staff to teach hundreds of thousands. However, as small teacher-to-students ratios can lead to substantial learning benefits (see e.g. Bloom [1]), scalable methods for providing students with more personalized instruction are of significant interest. Though peers and alumni are important resources (and provide an often important social component), both groups may have fewer insights into how to effectively teach the material. Therefore it would also be welcome to provide automated assistance of quality similar or exceeding an expert teacher. As a step towards this, we consider how to identify which dynamic

student features can be used to automatically predict the pedagogical activities selected by teachers.

Understanding the features used by experts to make instructional decisions is important for multiple reasons. First, the vast majority of computational student models focus on predicting the performance of a student. With few exceptions (e.g. [5]) little attention has been paid to creating student models explicitly designed to consider their impact on instruction. Indeed, that is the second key potential benefit: that armed with the knowledge of what features are used by expert teachers to select among pedagogical activities for a student, we could create automated methods to find highly effective instructional policies. Here we use the word policy to indicate a mapping from features about the student learning process to which next activity to give to the student. Reinforcement learning algorithms could help learn such policies from data, but to learn good policies typically requires an amount of data that scales prohibitively (from polynomially to exponentially) with the number of features used to describe the current state of the process. If expert teachers use only a small set of features, considering policies that only depend on such features could substantially reduce the amount of data required to find policies that meet or exceed human teacher performance.

One natural approach would be to directly ask teachers both what features they use when teaching. Unfortunately research from cognitive task analysis [2] suggests that experts may frequently fail to verbally specify all of the information they are using to make task decisions. To get at such information often requires additional effort beyond the interactions a teacher would normally have with students. Therefore it remains of use to directly infer features from existing data. A second issue is that the available data may be from good but imperfect teachers. Therefore we may not wish to be restricted to the specific policy provided by the teachers, and may wish to consider other policies that condition on the same feature set used by teachers.

Approach and Preliminary Results

Our approach follows recent exciting work on abstraction from demonstration [3]. The key idea of is to take demonstrations of a human user doing a task, infer the features of the task used by the human to make his or her decisions, and then perform reinforcement learning given those features in the task. Cobo et al. [3] demonstrated a small number of trajectories of humans playing simple video games was sufficient to identify a small set of features used by the players to make their decisions. The authors then used reinforcement learning with those features to efficiently learn a policy that well surpassed human performance.

We tried this approach in an education context, for a module about histograms. Our materials were developed as extensions of the Statway curriculum that was itself designed to help community college remedial math students.We selected this area given the enormous audience for remedial math, the rising workplace demand for statistics, and the documented challenge many students find when learning about histograms [4]. Our activities and assessments were designed to address some of the core learning objectives about histograms that are also often commonly misunderstood [4].

In our initial experiment we got data from 6 students. The student interacted directly with the computer for the test and activities. Each student took a pretest, then a human tutor hand selected an activity for the student, the student completed that activity, and then the human tutor selected a next activity based on what he or she observed about the student's performance on the tutor. This interaction was repeated until the human tutor thought the student had mastered the material (or no more practice activities were available) and then the student completed a post test. There was no additional assistance given by the human tutor to the student beyond specifying the next activity (or post test) to do, in order to best mimic the type of interaction a student would have with an automated (computerized) tutor.

Following the abstraction from demonstration approach, we then built a decision tree classifier to predict which activity type the human tutor would select for the student. The activities were classified into different learning objectives about histograms, as well as taking the post test. The classifier was provided with a large set of possibly relevant features about the teaching interaction, including the number of items completed about each learning objective, the type of the last done item, the amount of time spent so far, the amount of time spent on the last item, the student's pretest score, if the last item was completed correctly on the first try, and a number of other features. We then tried removing each feature and re-estimated the cross-validated accuracy at predicting the activity type selected by the teacher. We removed the feature that had the smallest negative cost on the overall accuracy, and then again tried removing each of the remaining features. We repeated this procedure until removing any further features resulted in a cross-validated performance accuracy of less than 95% of the maximal accuracy obtained during this removal process: we observed overfitting occurred when using the full feature set, and so the highest performance occurred with a subset of the features.

Though we started with over 50 features, we found that just four features were sufficient to achieve 64.5% accuracy student cross-validated accuracy (e.g. hold data from one student out as a test set and use the other students data to train, and repeat), which was within 95% of the best observed accuracy of 67%. The 4 features selected were: total time spent so far, the learning objective type of the activity the student just finished, how long of a time gap the student had between practicing the previous previous learning objective, and how many total times the student has practiced an activity about defining the center, spread and shape of a histogram.

Though this analysis is preliminary, it still reveals some interesting potential findings. One, the selected features do not depend on the student's performance (in terms of correctness or number of attempts). Indeed, prior work suggests that some teachers may mostly follow a fixed curriculum rather than adapting to a particular student (see summary in [6]). Second, a very small number of features could predict almost as well as the full set which is encouraging, since finding a good policy over this smaller set will generally require many less additional samples. This is important because each "sample" is a real student and may impact their learning.

Discussion

Our work on this is still in early stages and many open issues remain. One is that our classifier accuracy rates suggest that our human tutors may also be leveraging other features to make their instructional decisions (which would imply we should create a larger set of features to try to better capture human decisions) and/or that the different human tutors are either using different strategies or changing their strategies as they interact with more students. To better understand this, we may try to directly elicit feedback from teachers as to the features they consider when making pedagogical choices, as a way to augment features automatically identified.

Looking further, it could be interesting to see if other features beyond those used by human tutors are important to consider when selecting activities: indeed Van Lehn also suggests in his survey [6] that there is evidence that human teachers may be much less effective at building nuanced models of a student's state of knowledge and misconceptions, and such information may be relevant to producing the best policy. However, figuring out which features are important may again open up a much wider feature space which has implications for the speed of identifying a good policy. In our next steps we intend to first explore what instructional policies we can automatically learn using the limited features identified and how their performance compares to our human tutors.

Our proposed approach may be relevant to other domains and situations where a teacher is selecting among a relatively small amount of choices to best aid a student. However, if there are an enormous or unbounded space of options (such as if a teacher might ask a student to write a new essay on a completely new topic based on the student's previous essay) then it may be impossible or require an enormous amount of data to accurately predict a teacher's choices, and to infer the features relevant to those predictions.

To conclude, this is a step towards enabling human MOOC instructors demonstrate how to teach a few students directly, and use this data to identify features and learn a good policy that could be used to automatically provide personalized pathways to thousands of future students.

Acknowledgements

We thank Google for a focused research award and the study participants.

References

[1] Bloom, B. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher 13*, 6 (1984), 4–16.

[2] Clark, R. E., and Estes, F. Cognitive task analysis for training. *International Journal of Educational Research 25*, 5 (1996), 403–417.

[3] Cobo, L. C., Zang, P., Isbell Jr, C. L., and Thomaz, A. L. Automatic state abstraction from demonstration. In *IJCAI*, vol. 22 (2011).

[4] Kaplan, J. J., Gabrosek, J. G., Curtiss, P., and Malone, C. Investigating student understanding of histograms. *Journal of Statistics Education 22*, 2 (2014).

[5] Lee, J. I., and Brunskill, E. The impact on individualizing student models on necessary practice opportunities. In *EDM* (2012), 118–125.

[6] VanLehn, K. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist 46*, 4 (2011), 197–221.