Exploring Learner Model Differences Between Students

Michael Eagle[1], Albert Corbett[1], John Stamper[1], Bruce M. McLaren[1], Ryan Baker[3], Angela Wagner[1], Benjamin MacLaren[1], and Aaron Mitchell[2]

Human-Computer Interaction Institute[1] Department of Biological Sciences[2]
Carnegie Mellon University
[3]Graduate School of Education, University of Pennsylvania

{meagle, corbett, jstamper, bmclaren, awagner, maclaren, apm1}@andrew.cmu.edu, ryanbaker@gmail.com

Abstract. Bayesian Knowledge Tracing (BKT) has been employed successfully in intelligent learning environments to individualize curriculum sequencing and help messages. Standard BKT employs four parameters, which are estimated separately for individual knowledge components, but not for individual students. Studies have shown that individualizing the parameter estimates for students based on existing data logs improves goodness of fit and leads to substantially different practice recommendations. This study investigates how well BKT parameters in a tutor lesson can be individualized ahead of time, based on learners' prior activities, including reading text and completing prior tutor lessons. We find that directly applying best-fitting individualized parameter estimates from prior tutor lessons does not appreciably improve BKT goodness of fit for a later tutor lesson, but that individual differences in the later lesson can be effectively predicted from measures of learners' behaviors in reading text and in completing the prior tutor lessons.

Keywords: BKT, Genetics, Machine Learning, Student Modeling

1. Introduction

Learner models of domain knowledge have been successfully employed for decades in intelligent tutoring systems (ITS), to individualize both curriculum sequencing [1,2,3,4] and help messages [5,6]. Bayesian methods are frequently employed in ITSs to infer student knowledge from performance accuracy, as in the citations above, as well as in other types of learning environments [7], and Bayesian modeling systems have been shown to accurately predict students' tutor and/or posttest performance [1], [3], [8,9].

These models generally individualize modeling parameters for individual knowledge components (KCs, also referred to as skills) [10], but not for individual students. Several studies have shown that individualizing parameters for students, as well as for KCs, improves the quality of the models [1], [11, 12, 13]. These approaches to modeling individual differences among students have monitored student performance after the fact, in tutor logs that have been previously collected to derive individualized student parameters for the tutor module(s). While these efforts have proven successful, they don't achieve the goal of dynamic student modeling within an ITS, since estimating and using individualized parameters concurrently within a tutor lesson is quite difficult. In this paper we examine how well individual differences in student learning in a lesson of the Genetics Cognitive Tutor [8] can be predicted ahead of time from two types of prior online activities: reading instructional text and solving problems in prior tutor lessons. In the following sections we describe Knowledge Tracing, the on-line student activities, the predictors derived from students' reading and prior tutor activities, and our success in using these predictors to model individual differences in the tutor.

Bayesian Knowledge Tracing (BKT) estimates the probability that a student knows each of the knowledge components (KC) in a tutor lesson. It employs a two-state Bayesian learning model – at any time a student either knows or does not know a given KC – and employs four parameters, which are estimated separately for each KC. BKT is employed in Cognitive Tutors to implement *Cognitive Mastery*, in which the curriculum is individualized to afford each student just the number of practice opportunities needed to enable the student to "master" each of the KCs, which is generally operationalized as a 0.95 probability that the student has learned the KC.

Individual Differences. Knowledge Tracing and Cognitive Mastery generally employ best-fitting estimates of each of the four parameters for each individual KC but *not* for individual students. In this work, we incorporate individual differences among students into the model in the form of individual difference weights. Following Corbett and Anderson [1], four best-fitting weights are estimated for each student, one weight for each of the four parameter types, wL_0, wT, wG, wS.

In this paper we focus on four types of BKT models for the third lesson in a Genetics Cognitive Tutor curriculum on *genetic pathways analysis* to examine how well IDWs in a tutor lesson can be predicted from prior online activities. The four models are: (1) a standard BKT model (SBKT) with no individualization, (2) a model with best-fitting IDWs for lesson 3 (BFIDW-L3), (3) models with best-fitting IDWs from prior lessons, and (4) a model with predicted individual difference weights derived from earlier activities. We compare how much each of the three types of individualized models improves upon the non-individualized SBKT fit (1).

In an earlier study, Eagle et al [14] estimated individual difference weights for the first lesson in this curriculum before students began using the tutor lesson, based on six measures of students' reading performance and six measures of students' pretest performance. The predicted IDW model was about 40% as successful as the best-fitting IDW model. In a second study, Eagle et al [15] examined how well individual difference weights for the second lesson in the curriculum can be predicted from the same 12 reading and pretest measures, along with 10 measures derived from lesson 1: the 4 best-fitting IDWs from lesson 1, and 6 other measures of student performance in the lesson. In this study, the predicted IDW model was about 60% as successful as the best-fitting model. The predicted model improved the goodness of fit by 4.1%, reducing RMSE from 0.413 to 0.396, while the BFIDW model reduced RMSE by 6.8% 0.385. This study found that reading measures remained useful as predictors of IDWs across all these models, but that pretest measures became much less important as tutor-performance measures were incorporated into the models.

2. Discussion and Conclusions

We examine how well we can predict IDWs in lesson 3 with the same types of reading measures as in [14,15] along with an expanded set of tutor performance measures. This study examines methods for predicting individual difference weights for students in BKT learning parameters (intercept and rate) and performance (guess and slip) for the third lesson in a Cognitive Tutor curriculum. This is an important issue because integrating IDWs into an intelligent tutor lesson is easier if the IDWs can be assigned before the student starts working in the lesson. We evaluate the different estimated IDWs by examining how well they fit student performance in Lesson 3, compared to (1) standard SBKT with no IDWs, and (2) a model with best-fitting weights for Lesson 3.

We find that directly applying the best-fitting IDWs from either of two prior lessons in the curriculum, or from both lessons combined, does not appreciably improve goodness of fit for Lesson 3, compared to the SBKT model. In contrast, estimating lesson-3 IDWs from measures of students' prior reading performance, and performance in the two prior tutor lessons, is more successful; it is 60% as successful as the best-fitting Lesson-3 IDW model in improving the goodness of fit compared to the SBKT model.

Several secondary conclusions emerge. First, a prior study [15] obtained very similar success in predicting IDWs based on reading performance, pretest performance and a smaller set of tutor performance measures. This study demonstrates that IDWs can be successful predicted without including pretest measures. This is potentially important since pretests may not be available in online learning environments. Second, among reading time measures and a wide range of tutor performance measures, no category of measures emerged as an especially strong predictor of Lesson 3 IDWs; instead it appears that predictive success depends on a broad range of predictor variables. Finally, reading time measures prove to be useful predictors of students' problem-solving behaviors in a subsequent tutor lesson, including reading time measures for text on a topic unrelated to that tutor lesson. This suggests that the reading time measures may reflect knowledge-acquisition strategies, as well as any knowledge acquired.

Acknowledgements. This research was supported by the National Science Foundation via the grant "Knowing What Students Know: Using Education Data Mining to Predict Robust STEM Learning", award number DRL1420609.

3. References

1. Corbett, A.T., Anderson, J.R.: Knowledge tracing: Modeling the acquisition of procedural knowledge. User modeling and user-adapted interaction, 4, 253-278. (1995)
2. Mayo, M., Mitrovic, A. Optimising ITS behaviour with Bayesian networks and decision theory. International Journal of Artificial Intelligence in Education, 12, 124-153 (2001).

3. Shute, V.: Smart: Student Modeling Approach for Responsive Tutoring. User Modeling and User-Adapted Interaction, 5 (1), 1-44. (1995)
4. Ritter, S., Yudelson, M., Fancsali, S., & Berman, S.: How mastery learning works at scale. Proceedings of the Third ACM Conference on Learning @ Scale, 71-79. (2016)
5. Ganeshan, R., Johnson, L., Shaw, E., & Wood, B.: Tutoring diagnostic problem solving. In G. Gauthier, C. Frasson, K. VanLehn (eds.) ITS2000 Intelligent Tutoring Systems, LNCS vol. 1839, pp. 33-42. Springer, Heidelberg. (2000)
6. Conati, C., Gertner, A., & VanLehn, K. Using Bayesian networks to manage uncertainty in student modeling. User Modeling and User-Adapted Interaction, 12, 371-417. (2002)
7. Pardos, Z., Bergner, Y., Seaton, D. & Pritchard, D.: Adapting Bayesian Knowledge Tracing to a massive online course in edX. Proceedings of the Sixth International Conference on Educational Data Mining, 137-144. (2013).
8. Corbett, A., MacLaren, B., Kauffman, L., Wagner, A., & Jones, E. A.: Cognitive Tutor for Genetics Problem Solving: Learning Gains and Student Modeling. Journal of Educational Computing Research, 42 (2), 219-239. (2010)
9. Gong, Y., Beck, J., & Heffernan, N.: Comparing knowledge tracing and performance factor analysis by using multiple model fitting. In V. Aleven, J. Kay, J. Mostow (eds.) ITS2010 Intelligent Tutoring Systems. LNCS vol. 6094, pp. 35-44. Springer, Heidelberg. (2010)
10. Koedinger, K., Corbett, A. & Perfetti, C.: The Knowledge-Learning-Instruction (KLI) framework: Bridging the science-practice chasm to enhance robust student learning. Cognitive Science, 36 (5), 757-798. (2012)
11. Lee, J., & Brunskill, E.: The impact of individualizing student models on necessary practice opportunities. In: Yacef, K., Zaiane, O., Hershkovitz, A., Yudelson, M., Stamper, J. (eds.) EDM2012 Proceedings of the 5th International Conference on International Educational Data Mining Society, 118-125. (2012)
12. Pardos, Z. & Heffernan, N.: Modeling individualization in a Bayesian networks implementation of Knowledge Tracing. In De Bra, P., Kobsa, A., Chin, D. (eds.) UMAP2010 Proceedings of the 18th International Conference on User Modeling, Adaptation, and Personalization, LNCS vol. 6075, pp. 255-266. Springer, Heidelberg (2010)
13. Yudelson, M., Koedinger, K., & Gordon, G.: Individualized Bayesian knowledge tracing models. In: Lane, C., Yacef, K., Mostow, J., Pavlik, P. (eds.) AIED2013 Artificial Intelligence in Education, LNCS vol. 7926, pp. 171-180, Springer, Heidelberg. (2013)
14. Eagle M., Corbett, A., Stamper, J., McLaren, B., Wagner, A., MacLaren, B., & Mitchell, A.: Estimating individual differences for student modeling in intelligent tutors from reading and pretest data. In A. Micarelli, J. Stamper & K. Panourgia (Eds.) *Intelligent Tutoring Systems: 13th International Conference Proceedings*, 133-143. New York: Springer. (2016)
15. Eagle M., Corbett, A., Stamper, J., McLaren, B., Baker, R., Wagner, A., MacLaren, B., & Mitchell, A.: Predicting individual differences for learner modeling in intelligent tutors from previous learner activities. In L. Aroyo, S. D'Mello, J. Vassileva, J. Blustein (Eds.) *Proceedings UMAP 2016*, 55-63. New York: Association for Computing Machinery. (2016)