Towards sharing student models across learning systems

Ryan S. Baker[1], Bruce M. McLaren[2], Stephen Hutt[1], J. Elizabeth Richey[2], Elizabeth Rowe[3], Ma. Victoria Almeda[3], Michael Mogessie[2], and Juliana M.AL. Andres[1]

[1] Graduate School of Education, University of Pennsylvania
[2] Human-Computer Interaction Institute, Carnegie Mellon University
[3] EdGE at TERC
rybaker@upenn.edu

Abstract. Modern AIED systems develop sophisticated and multidimensional models of students. However, what is learned about students in one system—their skills, behaviors, and affect—is not carried over to other systems that could benefit students by using the information, potentially reducing both the effectiveness and efficiency of these systems. This challenge has been cited by a number of researchers as one of the most important for the field of AIED. In this paper, we discuss existing progress towards resolving this challenge, break down five sub-challenges, and propose how to address the sub-challenges.

Keywords: Student model sharing, AIED system integration, BLAP

1 Introduction

More and more students use learning technologies each year, a trend accelerated by COVID-19 [6, 14]. Schools often have students use several learning platforms, even within the same subject [4]. However, these learning technologies do not currently work together to support students. What one learning system determines about a student's skills and behaviors is generally not carried over to other learning systems, reducing both educational effectiveness and efficiency— if a student learns a topic several times, and multiple learning technologies need time to learn the same thing about a student.

This challenge, bringing together distinct learning technologies, has been repeatedly referred to as a key goal for learning technologies. Kay [11] argued for "lifelong user models…existing independently of any single application and controlled by the learner." It was also a key part of the fifth challenge, "Lifelong and Lifewide Learning," in the AI Grand Challenges proposed in [18]. Finally, it was one of six "Baker Learning Analytics Prizes" (BLAP) challenges [3]. This challenge, "Transferability: The Learning System Wall," was posed as not just transferring student information from learning system A to learning system B, but in improving a student model that is already successful in learning system B and using that improved model to change how system B supports students at runtime, improving learning outcomes. Intentionally conceived in a more specific fashion than previous challenges, the Transferability/Learning Wall challenge was designed to represent a stepping-stone to the visions proposed in [11] and [18]—while representing improvement for students in itself.

2 Prior Work

Although learning systems do not yet connect their student models, there has been some past work to integrate learning systems in other fashions. In this section, we review that literature and discuss why it remains a significant step to integrate two systems' student models in an actionable way.

One of the most well-known areas of relevant prior work is in standards for logging data and representing student models. The Caliper framework provides a large set of ways to represent data from a variety of types of learning activities seen in learning management systems but has less support for the types of activities seen in the more complex interactions in AIED systems [9]. xAPI attempts to offer support for representing and sharing the data from a broader range of learning activities [5].

Both these platforms can be used to integrate systems through connections such as the Learning Tools Interoperability (LTI) standards [10]. Still, the connections offered are very simple, such as specifying the correctness of an action. Neither framework provides functionality designed for sharing the type of complex student models used in modern AIED systems. One AIED project was able to develop a workaround for the LTI standard to support simple transfer of student model information between platforms [1], but the approach only worked in a single direction, for a single piece of information, and required a direct platform-to-platform connection. In another example, [7, 15] connected two learning environments into the same reporting system.

Other research has attempted to simulate a student model connection between different learning systems or activities, without actually connecting systems/activities to each other. [15,16] developed a mapping between the skills in two different learning systems and then tested it by administering paper tests to students and analyzed the degree of agreement between the skills (but solely from the paper test data). [8] analyzed whether student knowledge model estimates from one lesson in a Cognitive Tutor would improve knowledge estimation on later lessons in a Cognitive Tutor. [17] asked twenty subjects to use both a research paper recommendation system and a scientific talk recommendation system (with order randomized) and then analyzed whether the second system's recommendations would have been more accurate if the first system's data had been used. These studies established the feasibility and potential usefulness of connecting student models across learning systems, paving the way for the next step: actually making the connection between learning systems.

3 The Problem, Broken Down Into Its Constituent Parts

The problem of sharing student models between two learning systems in a meaningful way that improves student outcomes breaks down into five sub-challenges:

1) *Connection*: The two systems need to seamlessly and digitally connect to each other, whether via API, shared database, or another technical link, so that one system can use the other system's inferences to inform its behavior.

2) *Mapping Related Constructs*: The two systems need to have student models of similar or related constructs, each of sufficient accuracy to be practically useful, and a mapping between the constructs in each system is needed [16].
3) *Evidence Integration*: Each system needs to have a way to integrate evidence from the other system into its own estimates based on how strongly each system's evidence predicts behavior in the other system.
4) *A Good Reason*: There needs to be a practical reason for connecting the student models, e.g. the student model drives an automated intervention, or the student model helps with a teacher orchestration system.
5) *Demonstration of Benefit*: The intervention (whether automated or by teachers) driven by the shared student model needs to actually make a difference to student behavior and outcomes if properly delivered, but only for some students (i.e. a student model is actually needed; the intervention is not universally beneficial).

4 Potential Steps towards an Architecture and Student Model Integration Algorithms

There are many possible approaches to connecting and sharing information between two or more learning systems (sub-challenge 1, Connection): these approaches can generally be grouped into two categories, system-to-system direct connections, and server-mediated connections. System-to-system direct connections are likely the quickest approach but are also hard to scale more broadly. It will be difficult to develop an ecosystem of learning systems working in concert through direct connections between individual learning systems. Instead, it will be more scalable to build a single server to facilitate connections between many learning systems. This could be achieved by an external web service, shown in Figure 1, that different learning systems can post student model inferences to or request student model inferences from. This external service would also need to be able to securely maintain a mapping of student IDs in different learning systems, with some form of access control for school districts or learning system developers to authorize sharing between learning systems.

Assuming that the two platforms model similar or related constructs (sub-challenge 2, Mapping Related Constructs), and that these models drive practical interventions (sub-challenge 4, A Good Reason), the next step is to select an algorithm that each platform will use to integrate information from the other platform

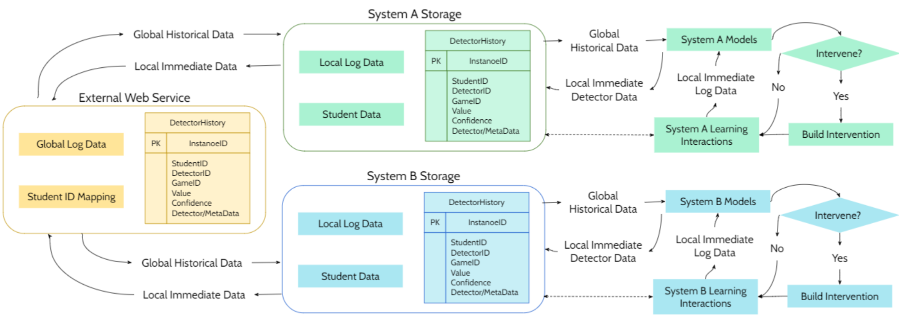

Fig. 1: A potential architecture for student model sharing

(sub-challenge 3, Evidence Integration), improving, replacing, or initializing the other system's estimates. Each system should take in the other system's evidence but make its own decision, rather than having a unified student model external to either system. This design choice keeps student model control local to each system—keeping system developers in control of their system's functioning. We propose investigating the following five approaches to information integration and selecting the most successful:

1) **System-weighted averaging**. Take each system's estimates, and average them together, weighting the other system's estimates lower than its own..
2) **System and evidence quantity weighted averaging.** Take each system's estimates and average them together. Each system weighs the other system's evidence in terms of the amount of evidence, penalized by a percentage due to the evidence not being from the local system.
3) **Performance Factors Analysis (PFA)** [13]. PFA is typically used in a single system. It computes a linear combination of weighted successes and failures for a skill so far (weights fit per skill) and then runs that combination through a logistic function to predict correctness on future items. PFA could be extended for a multi-system student model by fitting "successes" and "failures" for each system.
4) **Bayesian Network.** A Bayesian Network allows complex inter-relationships between skills [cf. 2, 12]. Both the current system's evidence and the other system's evidence can be integrated into a network, with the other system's evidence providing updates to the estimates of the current system's evidence.
5) **Deep Knowledge Tracing +** [19]. Deep Knowledge Tracing (DKT) can find complex relationships between multiple sources of evidence to predict future performance. DKT+ is an extension based on regularization that fixes problems with the original formulation (such as correct performance leading to predictions of worse performance and wild swings in proficiency estimates). The other system's evidence and the current system's evidence can be integrated into DKT+ to predict multiple student attributes or behaviors simultaneously.

Having integrated the two student models, the next step will be to test whether an intervention based on the integrated student model is beneficial for learners (sub-challenge 5, Demonstration of Benefit): beneficial only to students in need and better than an intervention from only a single system's data. One of the biggest areas of potential will be for "cold start" situations – where one system has evidence on student knowledge of a topic not yet encountered in the other learning system. There will also be potential around inferring constructs where considerable amounts of aggregate data are needed to draw a clear inference or where the behavior or state of interest only manifests occasionally.

5 Conclusion

In this article, we discuss the potential of sharing student models between learning systems. We frame this challenge in terms of five sub-challenges that need to be addressed in order to solve this challenge. We then offer an architecture to address a key sub-challenge and discuss algorithms that could potentially be used for student model integration. We encourage our AIED colleagues to join in solving this challenge.

References

1. Aleven, V., Sewall, J., Andres-Bray, J.M., Popescu, O., Sottilare, R., Long, R., Baker, R.: Towards deeper integration of intelligent tutoring systems: One-way student model sharing between GIFT and CTAT. In: Proceedings of the 7th Annual Generalized Intelligent Framework for Tutoring (GIFT) Users Symposium (2019)
2. Arroyo, I., & Woolf, B.P.: Inferring learning and attitudes from a Bayesian Network of log file data. In: Proceedings of the International Conference on Artificial Intelligence and Education, pp. 33–40 (2005)
3. Baker, R.S.: Challenges for the future of educational data mining: The Baker learning analytics prizes. Journal of Educational Data Mining, vol. 11 (1), pp. 1–17 (2019)
4. Baker, R.S., & Gowda, S.M.: The 2018 technology & learning insights report: Towards understanding app effectiveness and cost. San Francisco, CA: BrightBytes, Inc. (2018)
5. Berking, P., & Gallagher, S.: Choosing a learning management system. In: Advanced Distributed Learning (ADL) Co-Laboratories, pp. 40-62 (2013)
6. BrightBytes, Inc.: 2020 Remote Learning Survey Research Results. (2020). Retrieved February 12, 2021 from https://www.brightbytes.net/rls-research
7. Brusilovsky, P., Mitrovic, A., Sosnovsky, S., Mathews, M., Yudelson, M., Lee, D., & Zadorozhny, V. Database exploratorium: a semantically integrated adaptive educational system. InL *7th International Workshop on Ubiquitous User Modelling (UbiqUM 2009)*
8. Eagle, M., Corbett, A., Stamper, J., McLaren, B.M., Baker, R., Wagner, A., MacLaren, B., & Mitchell, A.: Predicting individual differences for learner modeling in intelligent tutors from previous learner activities. In: F. Cena, M. Desmarais, D. Dicheva, J. Zhang (Eds.), Proceedings of the 24th Conference on User Modeling, Adaptation and Personalization (UMAP 2016), pp 55-63. ACM, New York, NY (2016).
9. IMS Global: Caliper Analytics. (n.d., a). Retrieved January 4, 2021 from http://www.imsglobal.org/activity/caliper
10. IMS Global: Learning Tools Interoperability. (n.d., b). Retrieved January 4, 2021 from https://www.imsglobal.org/activity/learning-tools-interoperability
11. Kay, J.: Lifelong Learner Modeling for Lifelong Personalized Pervasive Learning. IEEE Transactions on Learning Technologies, Vol. 1(4), pp. 215–228 (2008)
12. Kim, Y. J., Almond, R. G., & Shute, V. J.: Applying evidence-centered design for the development of game-based assessments in physics playground. International Journal of Testing, vol. 16(2), pp. 142–163 (2016).
13. Pavlik Jr, P. I., Cen, H., & Koedinger, K. R.: Performance Factors Analysis--A New Alternative to Knowledge Tracing. In: Proceedings of the International Conference on Artificial Intelligence and Education. (2009)
14. Reich J., Buttimer, C.J., Fang, A., Hillaire, G., Hirsch, K., Larke, L., Littenberg-Tobias, J., Moussapour, R., Napier, A., Thompson, M., & Slama, R.: Remote Learning Guidance from State Education Agencies during the COVID-19

Pandemic: A First Look (2020). Unpublished manuscript. Retrieved February 12, 2021 from osf.io/k6zxy/.
15. Sosnovsky, S., Brusilovsky, P., Yudelson, M., Mitrovic, A., Mathews, M., & Kumar, A.: Semantic integration of adaptive educational systems. In: *Advances in ubiquitous user modelling*, 134-158 (2009)
16. Sosnovsky, S., Dolog, P., Henze, N., Brusilovsky, P., & Nejdl, W.: Translation of overlay models of student knowledge for relative domains based on domain ontology mapping. In: R. Luckin, K.R. Koedinger and J. Greer (eds.) Proceedings of the 13th International Conference on Artificial Intelligence in Education, IOS, pp. 289–296 (2007)
17. Wongchokprasitti, C., Peltonen, J., Ruotsalo, T., Bandyopadhyay, P., Jacucci, G., & Brusilovsky, P.: User model in a box: Cross-system user model transfer for resolving cold start problems. In: Proceedings of the International Conference on User Modeling, Adaptation, and Personalization, pp. 289-301 (2015)
18. Woolf, B.P., Lane, H.C., Chaudhri, V.K., & Kolodner, J.L.: AI grand challenges for education. AI Magazine, 34, 4, pp. 66–85 (2013)
19. Yeung, C. K., & Yeung, D. Y. Addressing two problems in deep knowledge tracing via prediction-consistent regularization. In: Proceedings of the Fifth Annual ACM Conference on Learning at Scale, pp. 1-10. (2018)