Integrating MOOCs and Intelligent Tutoring Systems: edX, GIFT, and CTAT

Vincent Aleven[1], Ryan Baker[2], Nathan Blomberg[3], Juan Miguel Andres[2], Jonathan Sewall[1], Yuan Wang[4], Octav Popescu[1]
[1]Carnegie Mellon University, [2]University of Pennsylvania, [3]Dignitas Technologies, LLC,
[4]Teachers College, Columbia University

INTRODUCTION

Massive Open Online Courses (MOOCs) are successful and widespread, but existing MOOCs have limited capacity to adapt to learners' individual characteristics. They often lack learner models that dynamically capture learner differences and can be used to drive adaptivity. Many have noted that MOOCs have both relatively low completion rates and relatively poor evidence for positive learning outcomes. It is likely that MOOCs would be more effective if they could personalize instruction based on learner characteristics, for example, learners' knowledge or personal interests. They might be more effective, further, if they would provide more learning-by-doing activities (Koedinger et al., 2015).

One path towards making MOOCs more adaptive, with a broader range of learning-by-doing activities, is to integrate existing adaptive learning technologies into MOOCs. In our ongoing project, we integrate the Generalized Intelligent Framework for Tutoring (GIFT) and the Cognitive Tutor Authoring Tools (CTAT) into the widely used edX MOOC platform. GIFT provides a framework and authoring tools for adaptive instruction (Goldberg & Hoffman, 2015; Goldberg, Hoffman, & Tarr, 2015; Sottilare, 2012). The CTAT tool suite can be used to build (among other things) example-tracing tutors, a type of intelligent tutoring system (ITS) that has proven to be robust, classroom ready, and effective (e.g., Aleven et al., 2016b). Like other ITSs, example-tracing tutors support learning-by-doing (practice on complex, recurrent problem types) with adaptive feedback and adaptive problem selection. ITSs have been shown to be very effective in enhancing student learning in a wide range of domains (Kulik & Fletcher, 2015; Ma, Adesope, Nesbit, & Liu, 2014; Steenbergen-Hu & Cooper, 2013; 2014; VanLehn, 2011) and can be authored, increasingly, with efficient and easy-to-learn authoring tools, such as CTAT.

The integration of GIFT and CTAT facilitates the addition of learning-by-doing activities to edX and enhances its adaptivity to learner differences. GIFT and CTAT have different roles to play in this integration, but both would enhance the adaptivity in MOOCs. The types of adaptivity offered by GIFT and CTAT are compatible but complementary, a prime reason why combining them has the potential to benefit learners. GIFT's Engine For Management of Adaptive Pedagogy (EMAP) module lets an author create policies for adaptively navigating Merrill's four quadrants (Goldberg et al., 2015). These quadrants characterize four types of instructional activities in a 2x2 grid defined by two dimensions: rules/instances and tell/elicit. EMAP policies for adaptively navigating these quadrants in an individualized manner can consider cognitive factors, metacognitive factors, and other factors. They are primarily outer-loop policies, meaning they take care of task selection (VanLehn, 2006). On the other hand, CTAT tutors offer complex problems with inner-loop adaptive support (VanLehn, 2006), meaning various forms of within-problem guidance, such as feedback on steps and next-step hints. CTAT, together with the TutorShop (Aleven et al., 2016b), a learning management system created in tandem with CTAT and geared specifically to ITS use, also offers outer-loop adaptive support, namely, Cognitive Mastery based on Bayesian Knowledge Tracing (Corbett, McLaughLin, & Scarpinatto, 2000), in which problem selection is individualized based on a student's knowledge. Our project will synergistically combine these different adaptive methods. We will demonstrate this integration and its potential for enhancing student outcomes in the context of the edX MOOC “Big Data in Education: Core Methods in Educational Data Mining” (BDEMOOC).

The principal software development challenges are, first, to integrate GIFT into edX and second, to integrate CTAT/TutorShop into GIFT. To meet these challenges, we take advantage of existing e-learning interoperability standards, namely, the Learning Tools Interoperability (LTI) and the Experience API (xAPI) standards (*IMS Global Learning Tools Interoperability™ Implementation Guide*, 2012; *xAPI Architecture Overview*, 2015). The project builds on our recent work, which has integrated intelligent tutors authored in CTAT into edX MOOCs (Aleven et al., 2015b; 2016a) using the LTI standard. We are currently working to create a first version of the adaptive BDEMOOC by integrating GIFT and CTAT/TutorShop separately into edX, with the course materials integrated into these platforms. This paper reports our work in progress. Specifically, we report on making GIFT an LTI Provider and on extending the content in the BDEMOOC, using the CTAT and GIFT authoring tools. This work addresses the first of the two challenges just mentioned; the second remains for future work.

Integrating GIFT into edX and other LTI Tool Consumers

To integrate GIFT into edX, as mentioned, we rely on the Learning Tools Interoperability (LTI) specification (*IMS Global Learning Tools Interoperability™ Implementation Guide*, *2012*). The LTI specification distinguishes application programming interfaces for learning management systems (LMSs) and for learning activity objects or “tools.” Here, the LMS is called the LTI Tool Consumer, while systems that provide the learning activities (via URLs) are LTI Tool Providers. In addition to edX, several other popular MOOC and e-learning platforms implement the LTI Tool Consumer interface, including Coursera, Canvas, Moodle and Blackboard.

A first step in our current project, therefore, is to make GIFT’s web delivery system an LTI Tool Provider. This first step achieves an initial configuration in which GIFT and CTAT are integrated separately into edX as LTI Tool Providers; this integration already enables enhanced adaptivity in edX, although, as discussed further below, it does not yet achieve the goal of integrating CTAT within GIFT. In a later phase of the project, we will also implement the LTI Tool Consumer protocol in GIFT, for integration with CTAT/TutorShop.

Integration details

We modified the GIFT framework to be an LTI Tool Provider using the LTI 1.1 specification. This GIFT extension makes it possible to embed a GIFT *course* into any LMS or e-learning platform that implements the LTI Consumer interface. Figure 1 shows the components in GIFT that were modified to allow a GIFT course to be run from a Tool Consumer, through an LTI launch request. Step numbers (N) in the following description refer to the circled numbers labeling data flows in the diagram.

To invoke an LTI Tool, from a page referring to LTI content, the Tool Consumer (i.e., the LMS, in this case, edX) (1) triggers an initial LTI launch request configured by the instructor in the Tool Consumer authoring software. The request data includes the reference to the GIFT course that is to be launched.

A new GIFT component (the LTI Tutor Servlet) was created to handle all incoming LTI launch requests. Each request is validated per the LTI specification and (2) the LTI user requesting access is initialized in the GIFT UMS database. Once the LTI user is initialized, the GIFT LTI Tutor Servlet (3) responds back to the Tool Consumer with an appropriate redirect URL. The redirect URL (4) directs the Tool Consumer to the (5) GIFT LTI Landing Page for the user. Once this page is displayed, GIFT (6) receives a request to load the GIFT course for the LTI user. The GIFT Dashboard Web Servlet is used to handle the load course operation. During this request to load the GIFT course, the LTI launch request details (7) are validated again for security purposes. The GIFT File Services API (8) is then used to load the GIFT course for the LTI user. A progress bar shows the user the progress of the course loading process.

After the course is loaded, the LTI user (9) requests to launch the GIFT course. The launch GIFT course request (10) is received by the GIFT Tutor User Interface (TUI) Web Servlet. The LTI launch request details are validated (11) another time for security purposes. Once the LTI launch request details are validated, GIFT (12) presents the LTI user with the started GIFT course in the web client. The LTI user then (13) is able to take the GIFT course until (14) it ends. Once the GIFT course ends, the LTI user is presented (15) with a final LTI Course Ended Page.

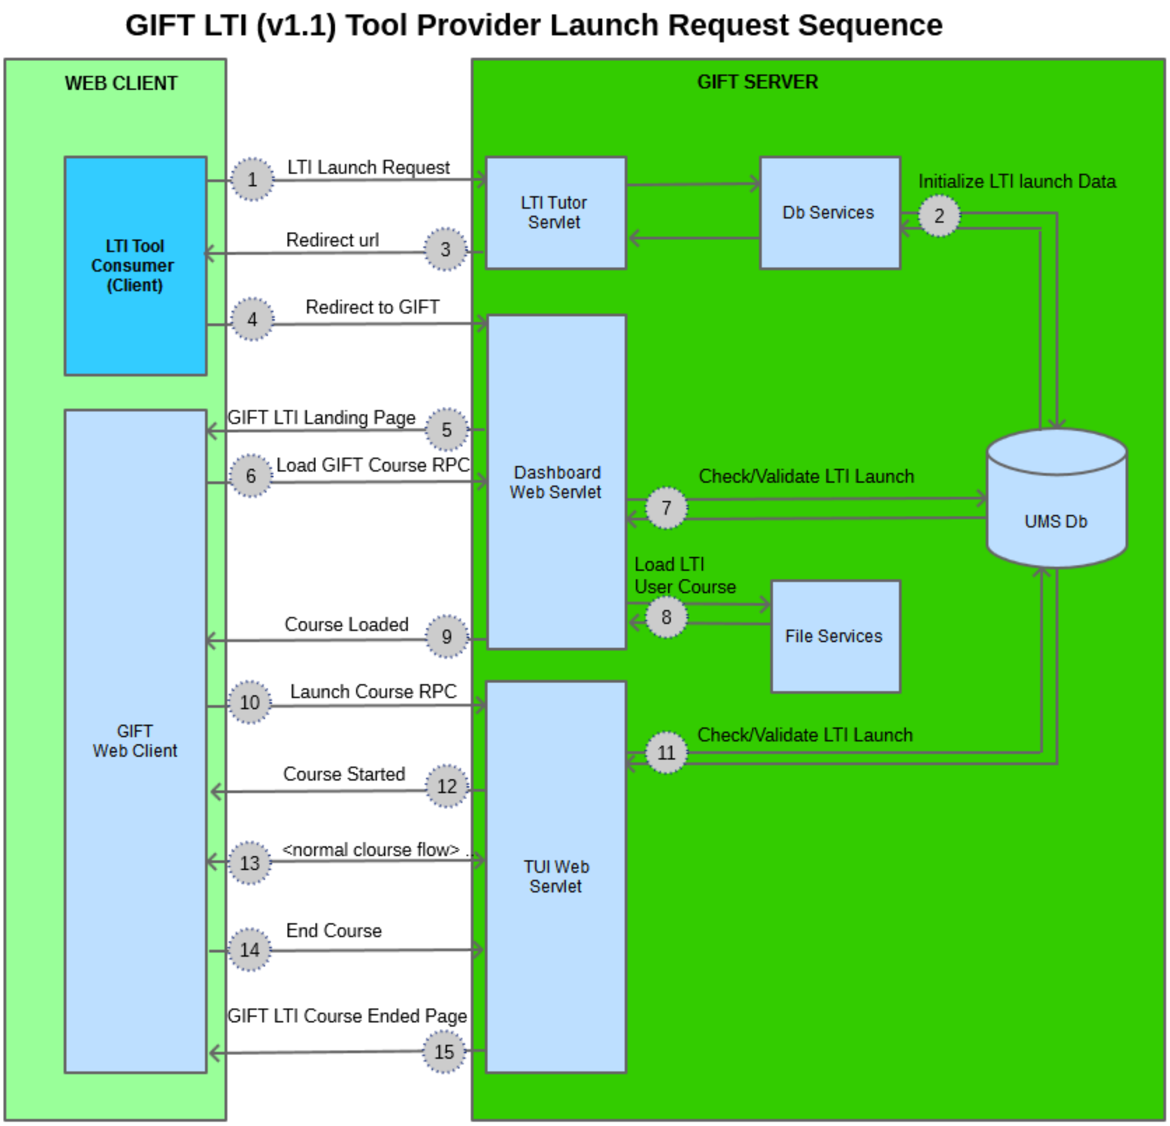

Figure 1. GIFT LTI Tool Provider launch request sequence

Challenges and next steps

We encountered a number of challenges during this effort. One challenge was anonymous users, since LTI users are not authenticated to GIFT as normal GIFT users are authenticated. To solve this challenge, the LTI Tutor Servlet now validates the LTI user and stores the LTI user information in a new database table.

A second challenge was a security issue in the way that GIFT courses were launched. Originally GIFT courses were referenced by a URL containing a human-readable path to the GIFT course file name. When this URL became the address for an LTI launch request, we found that it was exposed to students launching GIFT via LTI. Hence a malicious user could relatively easily access a different course by guessing and entering its file path instead. As a remedy, we replaced the file path in the URL with a universally unique identifier (UUID), which by itself provides no indication of the target GIFT file name or course.

Future integration efforts will provide the capability for GIFT to report a score to the LTI Tool Consumer, via a URL supplied as a parameter when the GIFT activity is invoked. In addition, a GIFT instructor will have the option to collect GIFT data on the users that have taken the GIFT course via LTI.

Adapting GIFT courses to run under edX

The unit of integration for GIFT into an LTI Consumer is an entire GIFT course. This makes it possible to embed any amount of GIFT content in a panel on a single page of an edX MOOC, even just a single exercise. The considerations behind this decision were at least partly technical: each entry to GIFT from edX requires authentication, and in GIFT a student authenticates himself or herself at the whole-course level, not for individual course modules. The result is that an existing GIFT course might need to be largely atomized into individual modules to be appropriate in the context of a MOOC. In our case, we may have whole GIFT courses consisting of just a single Adaptive Courseflow (i.e., EMAP) instance.

Future Work: GIFT as LTI Tool Consumer

In the final version of the MOOC we will create, we will use GIFT to select and deliver which activities and content the students receive, depending on their interest and knowledge, switching between CTAT-authored activities and GIFT-authored activities in the MOOC, and providing remedial conceptual instruction where appropriate. In this way, we combine the adaptive capabilities of GIFT (primarily, but not exclusively, for the outer-loop – the loop over tasks) and CTAT (primarily, but not exclusively, for the inner loop – the loop over steps within a task) in a more tightly integrated fashion than in the initial integration described above.

To integrate CTAT into GIFT, we will again we follow the LTI standards, but this time GIFT will become an LTI Tool Consumer. (As mentioned, CTAT/Tutorshop is already an LTI Provider.) That is, we leverage the LTI standards to enable any GIFT course to include *any* LTI-compliant learning tools, making available to GIFT not only CTAT tutors but also a host of learning activities beyond CTAT. This feature greatly expands the range of third-party-developed learning experiences available to GIFT courses, cuts costs, and eases development by re-using existing products. A sketch of the proposed final integration is in Figure 2. At the simplest level of interaction, CTAT (as it can already, under LTI v1.1) will return scores to GIFT as the student progresses through each exercise; GIFT could use these scores to decide whether or not to present remedial content. More detailed communication between CTAT and GIFT could be mediated by xAPI statements stored in and retrieved from a Learning Record Store (LRS).

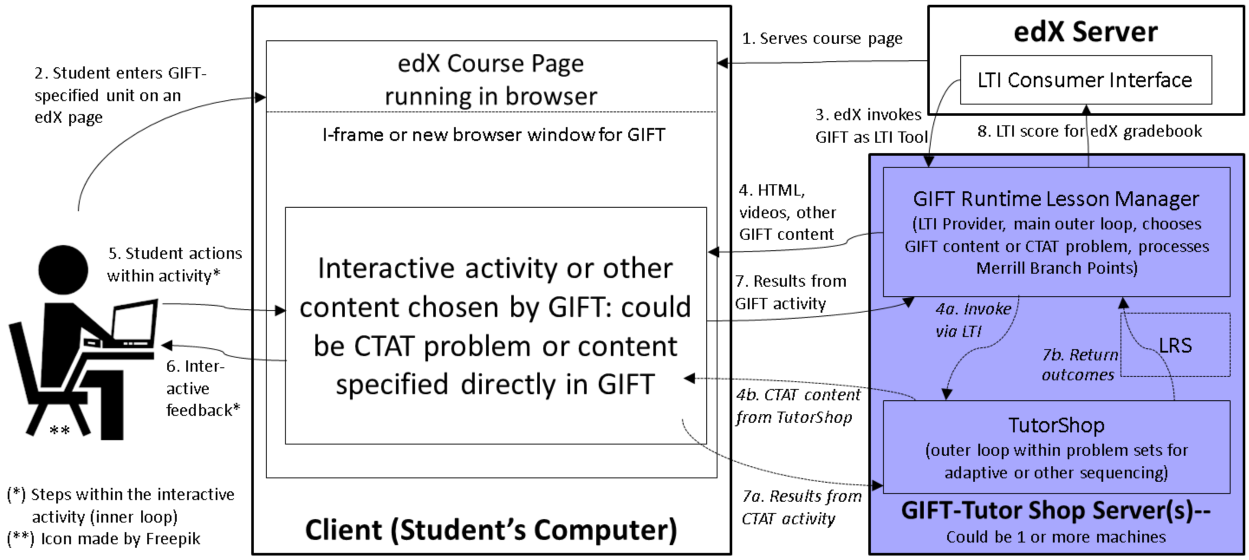

Figure 2: Final integration of GIFT and CTAT-TutorShop in edX, to be achieved in future work.

For example, take a student in the BDEMOOC (the edX course that will become our first adaptive MOOC) who has reported considerable interest in clustering, factor analysis, and bottom-up methods in general—all key topics in this course. This student's efforts will be directed towards these topics. In (say) week 2, rather than continuing further study of prediction models, as would be the default path through the course, the student will begin the study of clustering. Imagine that this student then successfully masters some of the key facts and concepts around clustering but afterwards skips straight to the practical "using the method" assignment in RapidMiner without viewing the videos on procedures for clustering. The student then struggles with the first CTAT assignment on clustering, specifically with respect to selecting how many clusters to use. With data from CTAT, GIFT recognizes that the student needs remedial support and that he or she has covered the facts and concepts. GIFT then uses this information to recommend that the student watch the video on procedures for selecting the number of clusters.

Adapting MOOC content for use with GIFT

The Big Data in Education course (BDEMOOC) has been offered in MOOC contexts across two iterations in the past. Its first iteration was on Coursera, and its second iteration was on edX. This section describes the ideas and the efforts involved in taking the existing curriculum and course materials and modifying them in order to utilize GIFT's and CTAT's adaptive capabilities.

The CTAT tutors embedded in the BDEMOOC ask students to come up with results in RapidMiner from the course's datasets and then query the students' understanding of these numbers. The hints in the tutors' questions, sometimes extending more than 10 levels deep, coach students to help generate their results; the tutors also provide feedback messages to explain and correct common misconceptions. An example of a CTAT tutor running in the BDEMOOC is shown in Figure 3 below.

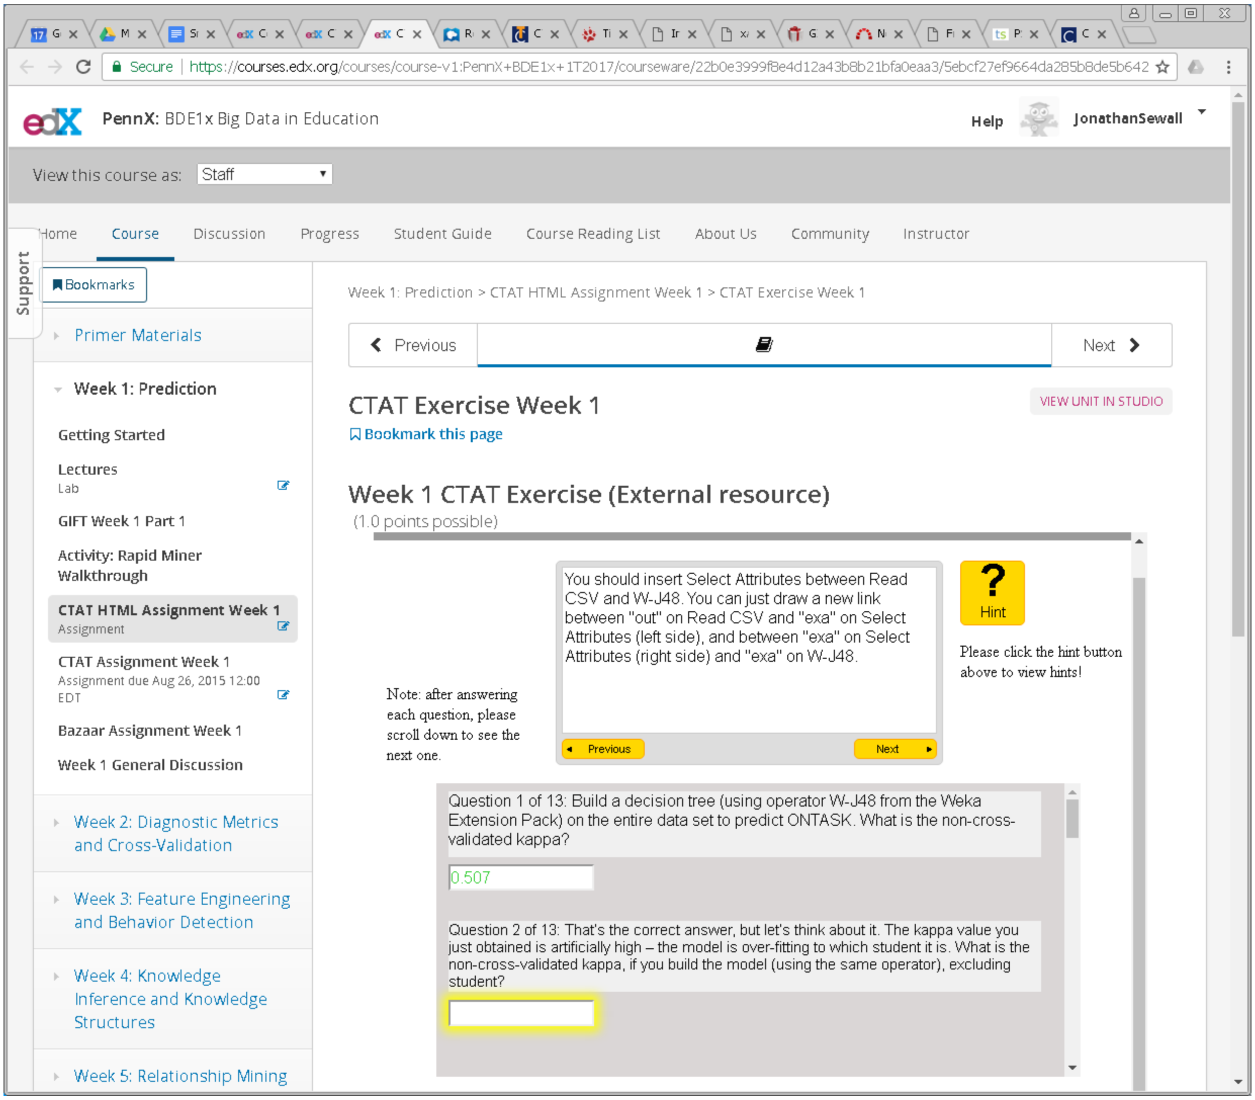

Figure 3: CTAT tutor running in the BDEMOOC. The student has asked for a hint on Question 2 and then requested additional help via the Next button on the Hint window. Because later questions could give the student clues to this answer, they are hidden until the student has answered this question correctly.

Identifying concepts

A major component of GIFT's learner model focuses on the student's competencies and how these are broken down into individual concepts. As such, the first step towards permitting GIFT to adaptively choose content within our course was to identify the individual concepts within the weekly modules. The previous iterations of BDEMOOC were comprised of 8 weekly modules, each having 4-8 video lectures. We identified approximately 200 concepts across all lectures. To give an idea of the granularity of these concepts, those from just the first half of week 1 are shown in Figure 4.

BDEMOOC2017 -- CDT matrix

fx | Topic

	A	B	C	D	E	F	G	H
1	Topic	Rules/Exposition	Examples	Recall	Practice	Conceptual	External Resources	
2	Glossary of things we'll cover in more detail later, but you need for {the walkthrough, ba1, ca1}							
3	Course Introduction	V1-1	V1-1	R1-1				
4	Three V's	V1-1 (full re-record)						
5	Diagram of topics in course					Wk 1		
6	What is Prediction	V1-2	V1-2	R1-2		CC1-4		
7	What is Regression	V1-2	V1-2	R1-3, R1-4		CC1-4		
8	Computing a Value for Regression	V1-2	V1-2		V1-2-quiz1			
9	Understanding Contribution of Variable to Model	V1-2			V1-2-quiz2			
10	Variable Transformations	V1-2	V1-2					
11	Benefits of Linear Regression	V1-2	n/a					
12	Risks of Interpretation with Multicolinearity	V1-2	V1-2					
13	Regression Trees	V1-2	V1-2	R1-7				
14	What is Classification	V1-3	V1-5	R1-5, R1-6	W1 (add sao ped	CC1-4		
15	Domain Specificity of Education	V1-3, V1-4	n/a					
16	Variables to Exclude in Classification	V1-3 (slides 49-5	V1-3		A1 (bde asgn.1 should exclude UNIQUEID; UNIQUEID bug; eliminat			
17	Step Regression	V1-3	V1-3		V1-3-quiz-1 (green), A1			
18	Logistic Regression	V1-3	V1-3, V1-5	R1-7, R1-8				
19	Logistic/Step Regression Limitations	V1-3	n/a			CC1-2		
20	Decision Trees	V1-3	V1-3	R1-8	W1, A1			
21	Decision Tree Benefits	V1-3	n/a			CC1-2		

Figure 4: Concepts from lectures in week 1 of Big Data in Education, mapped to Component Display Theory (plus conceptual activities that do not fit clearly in the theory). Specific videos are denoted V*week-video*, Recall items are denoted R*week-item*, Quizzes are denoted V*week-video*-quiz-*number*, Assignments are denoted A*week*, Walkthroughs are denoted W*week*, and Conceptual activities are denoted CC*week-prompt*.

We plan to create new content for BDEMOOC, as well as design and implement adaptive policies, primarily in GIFT, so the BDEMOOC activities can be selected and sequenced according to individual students' needs. In this initial iteration, we plan to take the first step towards making the course more adaptive through the modifications outlined below.

Course reorganization. Since none of the existing BDEMOOC course activities on edX currently have adaptive outer loop control, we will re-organize the course content in terms of Merrill's framework (on which GIFT's EMAP module is based) and embed it into GIFT and edX. These efforts will consist of three types of content creation: First, we will convert existing activities, so they can be used adaptively within GIFT or CTAT. Specifically, we will convert comprehension quizzes (currently existing within edX) to CTAT, divide existing content into four Merrill Quadrants, and remix/redivide videos for use in remediation and adaptive curriculum sequencing. Second, we will create new material for existing topic areas, to fully cover all four Merrill Quadrants (where reasonable; some topics may not lend themselves to practice, for example). This effort includes the creation of recall content. Third, we will create new material for new topic areas, to support deeper personalization of BDEMOOC to individual learners.

Adaption of content and creation of new content. In addition to reorganizing the course, we will bring more adaptivity into the activities within the course, taking advantage of the CTAT tools. Generally, CTAT-built tutors support multi-step problem-solving activities, with step-level support in the form of hint messages – offered on almost every step, to guide students through the thinking processes necessary to produce the correct answer, and buggy messages – adaptive guidance given when students provide a wrong answer that indicates a known misconception. Students working with a CTAT tutor can, at any point in the problem-solving process, voluntarily choose whether to access the hint messages and how many hint messages they would like to view. The initial versions of the CTAT tutors created for the BDEMOOC have been designed in a way that subsequent problem steps or questions will not appear until the student successfully completes the current step. To bring in more adaptivity for these activities, we plan to use the following methods: First, we will break existing tutor problems into smaller chunks and intersperse these smaller chunks between lecture videos. This way, learners will be presented with problem-solving steps relevant to the video lecture immediately preceding it. (By contrast, the current learning sequence has students going through all lecture videos during a week before attempting a single problem assignment at the end of the module.) Second, we will modify the hints in the CTAT-authored tutors to direct the learner to exact points in specific video lectures. In doing so, learners are encouraged to go back to watch the lecture video in order to revisit the learning content that is helpful toward answering the question correctly. The current hint framework allows the students to keep requesting hints until the final hint returns the answer to the question. Finally, in order to support deeper content-level personalization of BDEMOOC, we will create interactive activities for an additional fourteen topics1. These will supplement the current set of ten interactive activities built in CTAT for nine topics.

Running the adaptive BDEMOOC as real-world testbed

The new forms of adaptivity enabled by the GIFT/CTAT/edX integration will be piloted in BDEMOOC. This will provide a demonstration of the technical feasibility of the new forms of adaptivity created in the project. Specifically, in the upcoming years we plan to continue running BDEMOOC in edX at least three times. After each course run, we will iteratively improve and refine the course activities based on analysis of log data (as in the OLI Statistics course redesign based on KC analysis, Lovett, Meyer, & Thille, 2008). For data representing performance within CTAT, sent directly to DataShop, we will use the DataShop's learning curves and Performance Profiler to choose steps with the worst performance and revise those. Data on video watching will be analyzed in relation to assignment performance, in which students work with CTAT tutors, to see if students are watching the videos they need and appropriately using video resources to help them complete assignments. Forum data will be analyzed to determine how positive is students' sentiment to the modifications and what topics students continue to struggle with. In addition, we will gather questionnaire data to learn about how students experienced the adaptive MOOC. Attitudinal surveys will be analyzed to determine student attitudes towards the course.

CONCLUSIONS AND RECOMMENDATIONS FOR FUTURE RESEARCH

By integrating GIFT/CTAT/edX, we are creating a MOOC that synergistically combines three forms of adaptivity: GIFT EMAP adaptive outer-loop control, CTAT adaptive outer loop control, and CTAT adaptive inner loop functionality. We will take advantage of GIFT's adaptive method for navigating the four Merrill quadrants, selecting activities based on individual student characteristics. It is our goal that almost all educational activities will be embedded into GIFT and under GIFT EMAP control, with the exception of the course discussion forum. In the future, when GIFT invokes a CTAT problem set, it will be able to set various parameters that control how CTAT tutors operate within the given problem set. Within these problem sets, CTAT/Tutorshop will support adaptive problem selection, based on Bayesian Knowledge Tracing and Cognitive Mastery – this method has been shown to improve student learning, compared to giving all students the exact same set of problems (Corbett et al., 2000). CTAT will also provide adaptive support for multi-step problem-solving activities, with adaptive hints and feedback that can be sensitive to the student's path through the problem.

So far, we have achieved an initial integration in which GIFT and CTAT activities are embedded separately into edX. To achieve this integration, we made GIFT an LTI Provider. This initial integration

[1] Correlation Mining and Post-Hoc Testing, Causal Data Mining, Association Rule Mining, Sequential Pattern Mining, Differential Sequence Mining, Network Analysis, Epistemic Network Analysis, Clustering, Factor Analysis, Q-Matrices, Partial Order Knowledge Spaces, Bayesian Networks, Bayesian Knowledge Tracing, Performance Factors Analysis.

partially achieves the forms of adaptivity listed above. What is missing is that we do not yet support EMAP control over CTAT activities, which will further enhance the adaptivity (i.e., make it possible to schedule CTAT activities adaptively while also considering GIFT-based activities; this gives us two different authoring tools, each with their own strengths and limitations, to author activities in the different quadrants). In order to achieve GIFT EMAP control over CTAT tutor assignments, we will further extend GIFT so it implements the LTI Consumer protocol.

Our project will provide a practical framework for adaptivity in MOOCs, supported by proven authoring tools. The project will lay out a path for creating future GIFT/CTAT/EdX courses. Beyond creating such a MOOC, we will develop guidelines and a case study paper on how to effectively do this process going forward. Interestingly, the same LTI-based integration can be reused for embedding CTAT tutors or other LTI-complicant learning objects into GIFT courses (without edX), as well as for adding GIFT-based adaptivity (without CTAT) to edX courses or other LTI-compliant MOOC platforms, e-learning platforms, and LMSs.

The work will make a theoretical contribution by investigating how multiple instructional models can be combined: adaptive traversal of Merrill's quadrants, tutored problem solving by an ITS, and standard MOOC pedagogy focused on individual, self-regulated learning with a variety of online resources (video lectures, multiple-choice questions with automated feedback to test your understanding, practice problems with peer feedback, and online discussion with peers in forums). Because tutors and MOOCs were developed separately, their synergy is underexplored. Questions that we will investigate through this project include: What leverage is there in enabling CTAT to communicate its student model to GIFT, and how could GIFT use it to make better adaptive sequencing decisions? We will also study how the learner experience is impacted by this newer, more flexible pedagogy for MOOCs. We will survey learners to compare their attitudes towards the revised course to the attitudes reported in previous iterations of the course (both in terms of reported attitudes, and survey completion), and we will conduct sentiment analysis on discussion forum data to see how attitudes towards the course are impacted.

REFERENCES

Aleven, V., Sewall, J., Popescu, O., Ringenberg, M., van Velsen, M., & Demi, S. (2016). Embedding intelligent tutoring systems in MOOCs and e-learning platforms. In A. Micarelli, J. Stamper, & K. Panourgia (Eds.), *Proceedings of the 13th international conference on intelligent tutoring systems, ITS 2016* (pp. 409-415). Springer International Publishing.

Aleven, V., McLaren, B. M., Sewall, J., van Velsen, M., Popescu, O., Demi, S., . . . Koedinger, K. R. (2016b). Example-Tracing tutors: Intelligent tutor development for non-programmers. *International Journal of Artificial Intelligence in Education*, *26*(1), 224-269.

Aleven, V., Sewall, J., Popescu, O., Xhakaj, F., Chand, D., Baker, R. S., . . . Gasevic, D. (2015). The beginning of a beautiful friendship? Intelligent tutoring systems and MOOCs. In C. Conati, N. Heffernan, A. Mitrovic, & M. F. Verdejo (Eds.), *Proceedings of the 17th international conference on artificial intelligence in education, AIED 2015* (pp. 525-528). New York: Springer.

Corbett, A., McLaughlin, M., & Scarpinatto, K. C. (2000). Modeling student knowledge: Cognitive tutors in high school and college. *User Modeling and User-Adapted Interaction*, *10*, 81-108.

Goldberg, B., Hoffman, M. (2015). Adaptive Course Flow and Sequencing through the Engine for Management of Adaptive Pedagogy (EMAP). In *Workshop on Developing a Generalized Intelligent Framework for Tutoring (GIFT): Informing Design through a Community of Practice* (pp. 46-53), held in conjunction with the 17th International Conference on Artificial Intelligence in Education (AIED 2015).

Goldberg, B., Hoffman, M., & Tarr, R. (2015). Authoring instructional management logic in GIFT using the engine for management of adaptive pedagogy (EMAP). In R. Sottilare, A. Graesser, X. Hu, & K. Brawner (Eds.), *Design recommendations for adaptive intelligent tutoring systems* (Vol. III, Authoring Tools, Chap. 26, pp. 319-334). Orlando, FL: US Army Research Laboratory.

IMS Global Learning Tools Interoperability™ Implementation Guide (Final Version 1.1). (2012, March 13). Retrieved from https://www.imsglobal.org/specs/ltiv1p1/implementation-guide

Koedinger, K. R., Kim, J., Jia, J. Z., McLaughlin, E. A., & Bier, N. L. (2015). Learning is not a spectator sport: Doing is better than watching for learning from a MOOC. In G. Kiczales, D. M. Russell, & B. Woolf (Eds.), *Proceedings of the second ACM conference on Learning @ Scale* (pp. 111-120). New York: ACM.

Kulik, J. A., & Fletcher, J. D. (2015). Effectiveness of intelligent tutoring systems. *Review of Educational Research*, 0034654315581420.

Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*, *106*(4), 901. doi:10.1037/a0037123

Sottilare, R. (2012). Considerations in the development of an ontology for a generalized intelligent framework for tutoring. In I3M defense and homeland security simulation conference (DHSS 2012).

Steenbergen-Hu, S., & Cooper, H. (2013). A meta-analysis of the effectiveness of intelligent tutoring systems on K–12 students' mathematical learning. *Journal of Educational Psychology*, *105*(4), 970-987.

Steenbergen-Hu, S., & Cooper, H. (2014). A meta-analysis of the effectiveness of intelligent tutoring systems on college students' academic learning. *Journal of Educational Psychology*, *106*(2), 331-347.

VanLehn, K. (2006). The behavior of tutoring systems. *International Journal of Artificial Intelligence in Education*, *16*(3), 227-265.

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

xAPI Architecture Overview. (2015, October 23). Retrieved from https://www.adlnet.gov/adl-research/performance-tracking-analysis/experience-api/xapi-architecture-overview

ABOUT THE AUTHORS

Vincent Aleven *is an Associate Professor in Human-Computer Interaction at Carnegie Mellon. He has 25 years of experience in research and development of ITS based on cognitive theory and self-regulated learning theory, with a focus on K-12 mathematics. He and his colleagues have created the CTAT/Tutorshop ITS authoring environment, which non-programmers can use to create adaptive tutors. These widely-used tools improve the cost-effectiveness of tutor authoring by a factor of 4-8. Aleven has over 200 publications to his name. He and his colleagues and students have won 7 best paper awards at international conferences. He is or has been PI on 9 major research grants and co-PI on 10 others.*

Ryan Baker *is Associate Professor at the University of Pennsylvania, and Director of the Penn Center for Learning Analytics. Baker has developed models of student engagement in over a dozen online learning environments, and has led the development of an observational protocol and app for field observation of student engagement by over 150 researchers in 4 countries. His Massive Online Open Courses have served approximately 96,000 learners. He was the founding president of the International Educational Data Mining Society, and is currently serving as Associate Editor of two journals. Baker has co-authored published papers with over 250 colleagues.*

Nathan Blomberg *is a senior software engineer at Dignitas Technologies with over fifteen years of professional software development experience. Nathan holds a B.S. in management computer systems from the University of Wisconsin Whitewater. After graduation, Nathan spent over 12 years working with commercial software development in the video game industry helping create high profile AAA franchises including EverQuest, NCAA college football, and the Madden football games. He contributed to notable* *companies such as Sony Online Entertainment and Electronic Arts. Nathan joined Dignitas Technologies in 2014, where he is currently a developer for the GIFT and Marksmanship projects.*

Juan Miguel Andres *is a PhD student at the University of Pennsylvania, under the Teaching, Learning, and Teacher Education program. He is a member of the Penn Center for Learning Analytics, where he currently researches the generalizability of findings on student engagement and course completion in MOOCs.*

Yuan Wang *is a Doctoral Research Fellow in Cognitive Sciences at Columbia University Teachers College. Her current research focuses on Massive Online Open Courses (MOOCs) learner motivation and prediction of learner post-course development, with the goal of building frameworks toward operationalizing MOOC learner success. Specifically, her projects take a comprehensive approach by linking three sources of learner data: pre-course learner motivation, within-course learner engagement, and post-course development. Particularly, measurement of post-course development reflects both individual learner career development as well as advancement of communities of practice. Her recent publications are in Journal of Online Learning and Teaching and Journal of Learning Analytics.*

Octav Popescu *is a Senior Research Programmer/Analyst at CMU's Human-Computer Interaction Institute, working on the Cognitive Tutor Authoring Tools project.*

Jonathan Sewall *is a project director at CMU's Human-Computer Interaction Institute, working on the Cognitive Tutor Authoring Tools.*