# Beyond Log Files: Using Multi-Modal Data Streams Towards Data-Driven KC Model Improvement

Ran Liu
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
ranliu@cmu.edu

Jodi Davenport
WestEd
300 Lakeside Drive, 25th Floor
Oakland, CA 94612
jdavenp@wested.org

John Stamper
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
jstamper@cs.cmu.edu

## ABSTRACT
The increasing use of educational technologies in classrooms is producing vast amounts of process data that capture rich information about learning as it unfolds. The field of educational data mining has made great progress in using log data to build models that improve instruction and advance the science of learning. Thus far, however, the predictive and explanatory power of such models has often been limited to the actions that educational technologies can log. A major challenge in incorporating more contextually rich data streams into models of learning is collecting and integrating data from different sources and at different grain sizes. We present our methodological advances in automating the integration of log data with additional multi-modal (e.g., audio, screen video, webcam video) data streams. We also demonstrate several examples of how integrating multiple streams of data into the knowledge component (KC) model refinement process improves the predictive fit of student models and yields important pedagogical implications. This work represents an important advancement in facilitating the integration of rich qualitative details of students' learning contexts into the quantitative approaches characteristic of EDM research.

## Keywords
Multi-Modal Data Analytics, KC Model Improvement, Log Data, Structured Event Analysis of Multiple Streams (SEAMS)

## 1. INTRODUCTION
As student learning becomes increasingly conducted on computers and other digital devices, vast amounts of learning-related data are produced. Ideally, such data will provide a rich picture of student knowledge and behaviors (e.g., [8]). But predicting performance and generating pedagogical insight is limited, in the majority of cases, to the actions that digital systems can log. Computerized tutors are often used in a classroom context, and log data cannot capture all learning phenomena. A student working at a computer might be working independently with few outside influences. Alternatively, she might be in a lively classroom, with other students around her, talking and even offering suggestions. Data that capture the context surrounding educational technology use may add to and complement log data. In some cases, it may lead to critical insights.

Educational data mining analyses often omit additional contextual data for a number of reasons. Data on classroom context are difficult to collect. Data from different sources are often collected at different grain sizes, which are difficult to integrate. Here, we present work that extends educational data mining techniques to incorporate multiple modalities of data (computer log files, audio, screen videos, and webcam videos). We present methods we developed that help streamline both the collection of additional streams of data and the linkage across multiple streams. In two experiments, we then demonstrate the value of incorporating multi-modal, contextually rich data streams into established educational data mining techniques. In the first experiment, students use a chemistry virtual lab tutor and, in the second, students use an intelligent tutoring system to collaborate on fraction arithmetic.

Specifically, we extend methods of data-driven knowledge component (KC) model refinement [17] by incorporating, into the process, multiple streams of data spanning different modalities. We show that KC model improvements uniquely derived from these additional data beyond log files led to improved predictive models of student learning and behavior. These improved models of learning, in turn, can generate actionable knowledge for systems, students, teachers, and researchers.

## 2. BACKGROUND
### 2.1 Related Work
Recent work reflects a growing interest in multi-modal data analytics, particularly surrounding project-based, constructionist, and/or informal learning contexts [4, 18]. These efforts have focused on capturing divergent student strategies [4] and interactions that happen outside of a traditional computer tutor environment (e.g., with peers and with the physical environment [16]). Their primary goal is to make technologies supporting open-ended learning environments more scalable and to develop assessments appropriate for this type of learning.

Areas of research within the EDM community have also focused on collecting sources of data computer logs cannot capture to serve as "ground truth" labels in training log-data based detectors. These efforts have largely focused on modeling and detecting students' motivational and affective states [2, 8, 15]. For example, models can detect patterns of log data activity that precede affective states like confusion, frustration, and boredom. Physiological data may also be collected and used to develop models that can detect affective states from machine-readable signals, such as facial features, body movements, and electrodermal activity [14].

Outside of these pockets of the community, though, the majority of EDM research has focused exclusively on using log data to model learning. Building statistical models to predict step-level performance and data-driven KC model (or Q-matrix [3]) discovery are examples of major branches of EDM research that are typically limited to computer-logged data. In the present work, we demonstrate the value of expanding EDM research to include additional data streams that convey important contextual information about students' learning. We also present methodological advancements that improve the ease with which additional data streams can be collected and incorporated into educational data mining methods more broadly.

## 2.2 Data-Driven KC Model Improvement

Knowledge component models are an important basis for the instructional design of automated tutors and are important for accurate assessment of learning. Knowledge components (KCs) refer to units of knowledge representation (e.g., facts, concepts, or skills) that students need in order to solve problems. A KC Model maps a set of KCs mapped to a set of items or problem steps. Student models that are based on more accurate KC models produce better predictions of what a student knows based on their performance and, thus, result in better assessment and improved learning and instruction [11]. Cognitive Task Analysis is the traditional method for creating cognitive models of learning, but it requires subjective decisions and large amounts of human time and effort. Data-driven techniques of KC model discovery and refinement, when applied to large sets of educational data, can provide both more objectivity and reduce human effort.

A method developed by [17] leverages tools available in the PSLC DataShop [10] to identify potential improvements to a KC model in a data-driven manner. This method iterates through the following steps: (1) inspect learning curve visualizations and best-fitting statistical parameter estimates for the best existing KC model, (2) identify problematic KCs, (3) hypothesize changes to the KC model based on examining constituent problem content and applying domain expertise, and (4) re-fit the statistical model with the revised KC model and assess improvements in predictive accuracy. The premise for this method is that a hallmark of learning on a well-defined KC is a smooth learning curve that shows monotonic improvement in performance over time. KCs that lack these learning curve characteristics, but not because students are at ceiling performance, are likely to involve certain problem steps that require unlabeled difficulty factors or knowledge demands.

After a problematic KC is identified, its constituent problem steps must be examined in order to identify potential hidden difficulty factors. Thus far, this part of the process is limited to what computer log data. For example, a researcher might examine the error rates of the different constituent problem steps for the KC in question and the problem step names to gain clues about hidden difficulties. In the best-case scenario, the researcher might have access to the actual problem content for the dataset (as in [17]) and can apply domain knowledge to identify potential KC modifications. This step of content examination can be greatly enriched by additional streams of contextually rich data from the relevant moments of learning. To this end, we present a method of integrating streams of contextual audio and video data into the KC model refinement process. We show that such integration leads to insights that would not be derived by solely analyzing log data or curriculum content in isolation. We present several examples of how these insights lead to quantitative KC model improvements that improve the overall fit of student models to the data.

# 3. METHODS

We developed a method of semi-automatically extracting epochs, across multi-modal data streams, associated with the moments during which students engage with a particular KC of interest. This allows the content reviewer, after identifying a candidate KC, to not only view the curriculum content associated with a given KC but also to experience students engaging with that curriculum content through multiple modalities.

There are many ways to collect additional streams of contextually rich data (e.g., using video cameras, external microphones, eye-trackers, and sensors). We focused on a method that minimizes both deployment effort and interference with students' usage of educational technology to increase the likelihood that researchers would consider collecting, analyzing, and sharing such data. In the following experiments, we used Camtasia to simultaneously capture audio recordings, screen videos, and webcam videos of the students. Camtasia can be run in the background to collect all of these streams of data while a student engages with educational software. We installed Camtasia to all classroom laptops in advance of the two studies. On each day of the studies, we opened Camtasia and prepared recording settings before each class period so that all students needed to do was click a red "Record" button prior to logging into the tutors. At the end, students were led through a simple sequence of steps to ensure that their recordings were saved and named properly for easy post-hoc identification.

All recordings (audio, screen video and webcam video) for a single session are initially saved in a Camtasia-specific file format. We used the batch processing function to import and convert the original files to MP4 files that contained all data streams merged. We used timestamp information within the log files to map segments of log data to the appropriate corresponding multi-modal video stream. This step required human input, as Camtasia does not automatically log the system time (at millisecond level) that marks the start of the video recording. For each video file, someone must identify the offset between the beginning of the video and the time of some event in the log file. This offset can then be used to automatically align all remaining events between the log file and the corresponding video files.

We developed a tool called Structured Event Analysis of Multiple Streams (SEAMS) that builds upon the moviepy Python package in conjunction with the FFMPEG multimedia framework to automatically extract video epochs associated with specific events in the log data. The tool allows the user to indicate any event type that can be identified by labels within the log data and generates a folder of video clips that contain all epochs of the merged data streams pertaining to the particular event of interest (in this case, a specific KC at the specific opportunity count). With the relevant epochs grouped together in a manner that allows for quick and effortless analysis by a human examiner, it becomes much easier to quickly view multi-modal data streams to identify hidden knowledge demands towards KC refinement.

We applied our methods to examine the contributions of additional multi-modal data streams on KC model refinement across two classroom experiments. One experiment engaged students in a Chemistry Virtual Lab tutor for which we collected both screen videos and webcam data of learners' facial expressions in addition to traditional log data. The other experiment engaged students in a Collaborative (partner-based) Fraction tutor, and we collected screen videos and audio recordings of students' collaborative dialogue. Due to processor limitations of the school laptops that were available for the Collaborative Fraction tutor experiment, we were not able to collect webcam data. Using the data from both of these studies, we illustrate the application of our methods to leverage the additional multi-modal data streams to improve upon existing KC models. These KC model improvements, in turn, yielded insights about how to improve instruction within the respective tutors.

## 4. EXPERIMENTS

### 4.1 Chemistry

ChemVLab+ (chemvlab.org) provides a set of high school chemistry activities designed to build conceptual understanding and inquiry KCs [6]. In each activity, students work through a series of tasks to solve an authentic problem and receive immediate, individualized tutoring. As students work, teachers are able to track student progress throughout the activity and attend to students that may be lagging behind. Upon completion of the activities, students receive a report of their proficiency on targeted KCs, and teachers can view summary reports that show areas of mastery or difficulty for their students. In the current study, students completed four modules: PowerAde: Using Sports Drinks to Explore Concentration and Dilution, The Factory: Using a City Water System to Explore Dilution, Gravimetric Analysis, and Bioremediation of Oil Spills.

#### *4.1.1 Participants*

Participants were 59 students at a high school in the greater Pittsburgh area enrolled in honors chemistry classes. They participated in four Stoichiometry modules of the ChemVLab+ educational tutor. They completed these modules across four 50-minute class periods spread over the course of 3 weeks. We collected, using Camtasia, audio recordings and screen video captures for 58 students and webcam recordings of facial expressions for a subset of 25 students who were comfortable with their face being recorded during tutor use.

#### *4.1.2 Results*

The newly developed methods facilitated the identification of the way in which a problematic KC needed to be split as well as technical issues that impacted student learning. First, following methods described in [17], we identified a knowledge component called *Concentration* that seemed to have uncharacteristically high error rates on later practice opportunities (Figure 1). This KC represents understanding that the measure of concentration is the amount of substance (e.g., a sports drink powder) in a volume of substrate (e.g., water). It also represents being able to read, report, and compare concentrations of solutions.

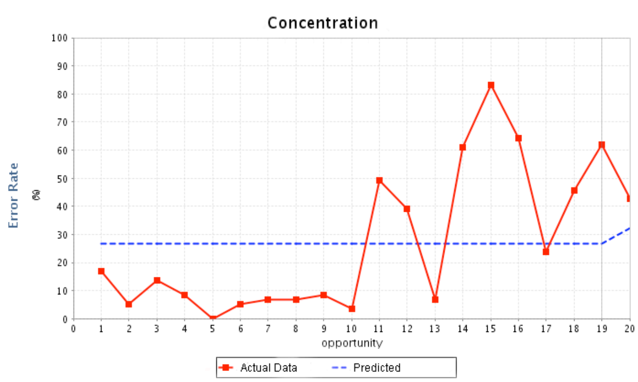


**Figure 1. Aggregate learning curve for the *Concentration* KC as originally defined by the ChemVLab+ tutor.**

We then used the methods described in Section 3 to automatically extract the screen and webcam videos of all epochs of students engaging with the Concentration KC on their 11th, 12th, 14th, 15th, 16th, and 19th practice opportunities. These were the opportunities on which the KC learning curve had unusually high error rates.

Qualitative analyses of these video stream epochs revealed that students were particularly confused by problems that involved dilution in conjunction with concentration, particularly when a dilution ratio or "factor" is involved. Students demonstrated this confusion as they responded to prompts such as 'Create a 1:2 dilution of the reported sample' or 'Add water to the sample until the concentration is diluted by a factor of 2'. The correct solution requires students to know that the amount of substance (e.g., the powder) takes up negligible volume, so to dilute the powder by 2x, the total amount of water needs to be doubled. Students demonstrated shallow knowledge by responding to prompts like these by adding two parts water to one part solution rather than adding one part water to one part solution, which halves the concentration. In another example, prompt 'Dilute this sample by a ratio of 6:1' student tended to add six parts water to one part of solution (making the resulting amount of powder to volume 1:7), part rather than adding five parts of water to one part of solution (making the resulting amount of powder to volume 1:6).

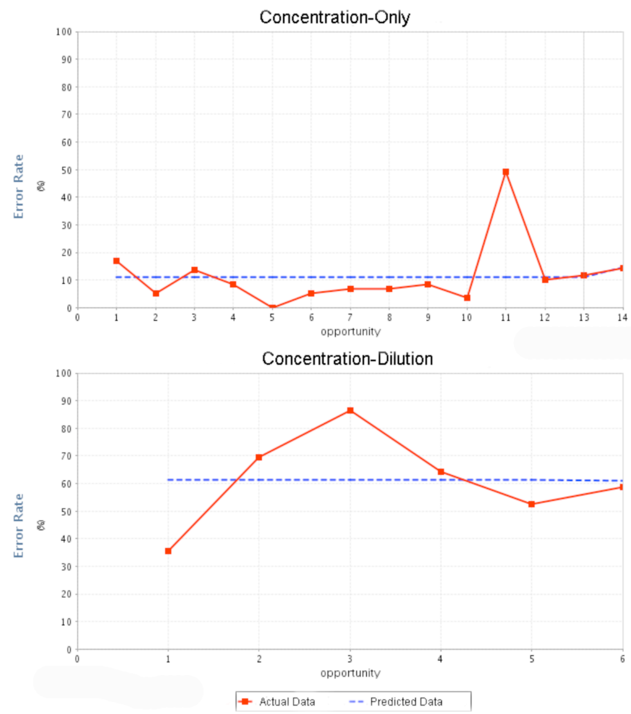


**Figure 2. Aggregate learning curves for the two new KCs, *Concentration-Only* and *Concentration-Dilution,* resulting from the KC model refinement process.**

Based on this insight, we split the Concentration KC into cases where the problem step required a conceptual understanding of dilution ratios/factors (Concentration-Dilution) and cases where it did not (Concentration-Only). The learning curves for the resulting two KCs are shown in Figure 2. The curves are much smoother than the original learning curve, with the exception of a particular opportunity count with unusually high error rate in the resulting 'Concentration-Only' KC at practice opportunity 11.

To further examine this unusual blip, we re-applied our method to automatically extract screen and webcam videos of all epochs of the 11th opportunity to practice the Concentration-Only KC. We noticed that the majority of problem steps experienced by students on this opportunity count were from a particular screen in the tutor in which the problem text was cut off in the interface. This resulted in students being confused about what they should be doing on this problem. Guessing the answer incorrectly was a common first attempt, as was clicking a hint button. Since the problem text was fine when viewed on research computers, it did not appear to be a problem with the educational software itself. We hypothesize that the problem may have been due to a unique interaction between the software and the resolution of the computers that students were working on. This is a reality of educational technology deployment in classrooms, and it would have been impossible to know from strictly the log data file or even problem content records that this was the source of students' struggle. If we had only accessed the recorded (idealized) version of the problem content, we may have incorrectly attributed the high error rate on this problem step to intrinsic content present within the problem. After separating these problem steps out from the 'Concentration-Only' KC, the resulting learning curve was much smoother, with an overall low error rate (Figure 3).

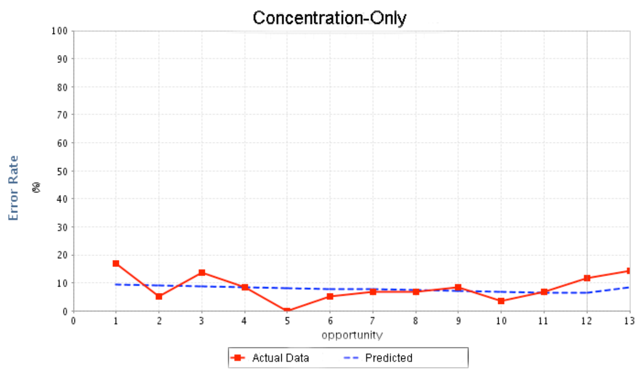


**Figure 3. Resulting *Concentration-Only* learning curve, after separating out the problem step in which students experienced a technical difficulty during deployment.**

Student model predictive fit metrics are shown in Table 1 for the different KC models when used in conjunction with the Additive Factors Model [5] and reveal an improvement in predictive fit across all metrics (AIC, BIC, and 10-fold cross validation) after splitting the original Concentration KC based on our qualitative analysis of student behavior during epochs of that KC (Row 2). Further improvements in predictive fit across all metrics were observed after we separated out the problem step that contained missing problem text during implementation (Row 3).

**Table 1. Student model fit metrics comparing different models resulting from the KC model refinement process.**

| | AIC | BIC | Cross Validation RMSE |
|---|---|---|---|
| Original KC model | 6694.58 | 7196.59 | 0.3859 |
| 'Concentration'-Split KC Model | 6388.35 | 6904.12 | 0.3838 |
| 'Concentration'-Split KC Model with text-error problem step separated | 6318.95 | 6848.47 | 0.3819 |

Both of these KC model refinements, each of which resulted in a substantive and consistent improvement in predictive accuracy when used by the Additive Factors Model, were uniquely dependent on qualitative analyses of the video data we had collected using Camtasia. Although it may have been possible to recognize that the concept of dilution ratios was an additional difficulty factor by purely accessing problem content, there were many other differences between the high error-rate problem steps and the low error-rate problem steps that constituted the original Concentration KC. For example, many of the higher error rate problem steps were part of a different activity (Activity 2, The Factory) than the lower error rate problem steps were (Activity 1, Powerade). Only by observing the students specifically exhibiting actions suggestive of possessing a shallow understanding of dilution ratios (via Camtasia screen videos) and affective states resembling frustration (via webcam videos) were we able to quickly identify the true hidden difficulty factor. Another benefit of this insight, perhaps even more significant than generating a better fitting KC model, is that there are clear implications for instructional redesign. That is, future iterations of the ChemVLab+ tutor might include instruction that more directly targets the misconceptions that students seem to have about the relationship between dilution ratios and existing solutions.

Discovering the high-error-rate problem step in which text was cut off would not have been possible without viewing the real context in which students experienced the problem. Since it was not a general problem with the ChemVLab+ tutor but, rather, an idiosyncrasy in that problem's display on the technology used in the classroom, the Camtasia screen videos were critical in correctly attributing the source of these errors.

## 4.2 Collaborative Fraction Tutor

The collaborative fraction tutor is online software developed by researchers at Carnegie Mellon University that helps students become better at understanding and working fractions. The tutor was created using Cognitive Tutor Authoring Tools, which allow for rapid development and easy deployment of intelligent tutors [1]. This particular fraction tutor supports collaboration between partners in order to learn fraction-solving KCs such as addition, subtraction, comparing fractions to determine which is larger or smaller, finding the least common denominator, and finding equivalent fractions. In the tutor, each student in a pair can control only part of the screen, so both partners must to work together in order to finish the problem. One student cannot do the whole thing him or herself. Students work at the same time and can talk about what they are doing, ask for help from their partner, and generally collaborate to get the correct answer.

### *4.2.1 Participants*

Participants were 26 fifth grade students at a middle school in the greater Pittsburgh area enrolled in an advanced math class. Students participated across five 45-minute class periods on consecutive days within a week. On the first and last days, students took a computerized pre- and post-test, respectively. They engaged in the Collaborative Fraction Tutor during the three consecutive days between the pre- and post-test days. Students spent half of each class period working individually and half collaborating with a partner. Students were paired with the same person for all partner activities throughout the experiment. We also collected audio and screen video captures for all students working both individually and in pairs on the three tutor use days.

### *4.2.2 Results*

The newly developed methods facilitated the identification of KCs that needed to be split. First, as in [17], we identified a knowledge component called *LCD_procedural* that was noisy, in particular due to an uncharacteristically high error rate on the 5$^{th}$ practice opportunity (Figure 4). We then used the methods described in Section 3 to automatically extract the combined audio and screen videos of all epochs of students engaging in their 5$^{th}$ opportunity of the LCD_procedural KC. Based on qualitative analyses of the video and audio streams, it was clear that the most common mistake that students were making on those practice opportunities was multiplying the two denominators but failing to reduce the product to find the least common multiple. This was particularly apparent in students' collaborative dialogue following their incorrect first attempts. Students often verbalized the realization that there must be a smaller common multiple. This verbalization did not occur on problems in which the product of denominators happened to be the correct solution. This suggests that there was a separate learning curve for the additional difficulty factor of cases where finding the least common denominator required reducing the product of the two original fractions' denominators to find a smaller common multiple. Based on this, we split the LCD_procedural KC into cases where the LCD required reducing from the product of denominators (LCD_procedural_REDUCE) and cases where it simply was the product of denominators (LCD_procedural_PRODUCT). The resulting learning curves (Figure 5) are much smoother than the original learning curve.

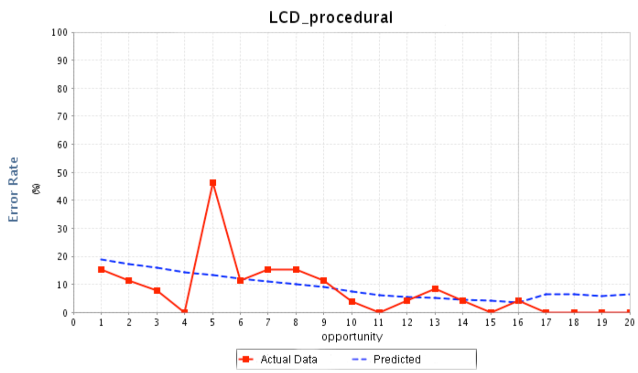


**Figure 4. Aggregate learning curve for the *LCD_procedural* KC as originally defined by the Collaborative Fraction tutor.**

The student model predictive fit metrics (Table 2) for the different KC models, when used in conjunction with the Additive Factors Model, reveal a substantial improvement in predictive fit across all metrics (AIC, BIC, and 10-fold cross validation) after splitting the original LCD_procedural KC based on our qualitative analysis of student behavior during epochs of that KC.

Through the audio-video segments, we observed students make denominator-product-based errors on their incorrect first attempts and realize they needed to find a smaller common multiple on certain problem steps. This greatly streamlined our identification of the hidden difficulty factor. As a result, we were able to quickly identify the appropriate KC split that led to much smoother learning curves and a better fitting student model.

This discovery also has important instructional implications: for example, the tutor might incorporate a bug message specific to students' inputting the product of the two denominators when the answer is a smaller multiple (i.e., "Can you find a smaller number that divides both denominators?"). A student model based on the revised KC model (with 'LCD_procedural' split into two separate KCs) would also result in students receiving more practice on problems in which the correct answer is a smaller multiple than the product of the two denominators. These instructional changes, resulting from the audio dialogue and video driven insights, will give students better support to overcome this difficulty.

**Table 2. Student model fit metrics compared between the original KC model and the improved KC model resulting from multi-modal data stream driven refinement process**

| | AIC | BIC | Cross Validation RMSE |
|---|---|---|---|
| Original KC model | 3497.6 | 4156.3 | 0.2738 |
| 'LCD_procedural' split KC model | 3462.2 | 4134.5 | 0.2734 |

## 5. DISCUSSION & FUTURE WORK

The vast majority of EDM research, especially research focused on predicting student performance and generating pedagogical insights, is limited to models based on computer-logged data. A recognized issue within the EDM community is that log data cannot capture all learning phenomena; it can miss important details of both learning processes and the learning context. Recent advances in DataShop [10] allow researchers to connect problem names in log data to screenshots of problem content and encourage inclusion of contextual details in custom fields of log data. Clearly, however, there are still instances where a better understanding of the implementation environment and students' experience working through certain problem steps is needed, as demonstrated here.

The main contributions of this work are (1) developing methodological advancements (e.g., the SEAMS tool) that facilitate the ease with which EDM researchers can incorporate context-rich data streams into quantitative modeling techniques, and (2) demonstrating the utility of doing so. Using a top-down,

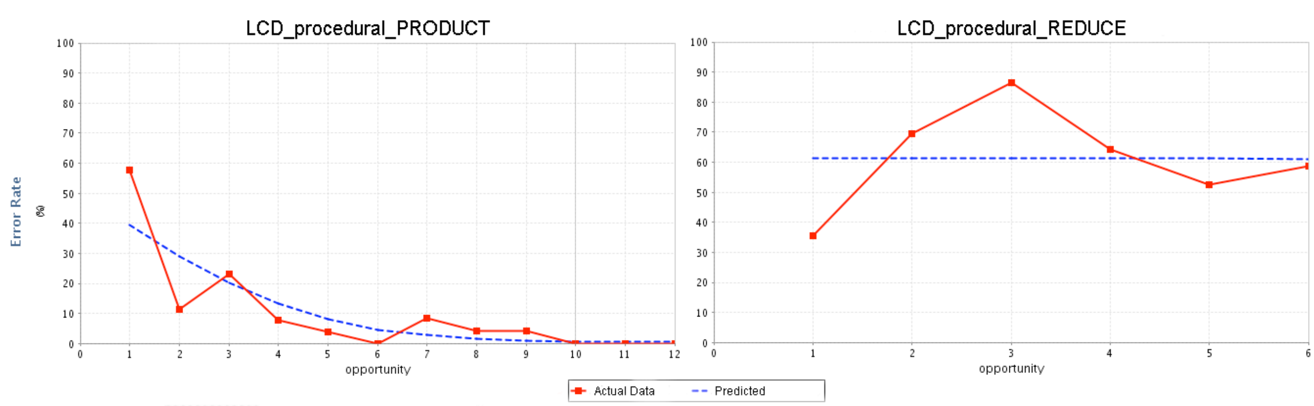


**Figure 5. Aggregate learning curves for the two new KCs, *LCD_procedural_PRODUCT* and *LCD_procedural_REDUCE*, resulting from our KC model refinement process5.**

KC visualization driven method, we show that valuable qualitative insights can be obtained from targeted segments of audio and video data even without fully "coding" all of the multiple streams. We also show that these qualitative insights lead to quantitative model fit improvements and actionable pedagogical implications.

There are many promising areas for future work based on the methods we have developed here. The present work has focused on refining an existing KC model. Educational data does not always come with an existing expert-labeled KC model, and there have been recent efforts to automatically generate, or discover, KC models [9, 12, 13]. One concern about fully machine-discovered models is their interpretability. The ability to view contextually-rich audio and video segments corresponding to machine-discovered KCs will facilitate the interpretation of these KCs and, in turn, help researchers refine their methods to yield more interpretable or cognitively plausible KC models.

Another interesting issue that contextually-rich streams of data are uniquely suited to address is the attribution of pauses of activity in the log data. A pause in the data because a student is off-task has very different implications than a pause because the student is actively help-seeking outside of the educational technology interface. Being able to use detailed information about students' learning context can help produce correct interpretations of log data activity and, in turn, more robust student models.

Finally, one of the interesting data streams we collected in the Chemistry dataset was student-facing webcam video. Aside from noticing the moments during which students seemed frustrated in the Chemistry tutor due to confused about dilution ratios, we have not yet fully explored the extent to which the webcam data could be used to improve KC models and student models. There is rich potential for our methods to facilitate connections between the cognitive (e.g., knowledge component modeling) and the affective [2, 8] branches of EDM research.

## 6. ACKNOWLEDGMENTS

We thank Jacklyn Powers and Jenny Olsen for help in collecting the Chemistry and Math tutor data used in our experiments here. This research was supported by the National Science Foundation (Grants #DRL-1418072, PI Davenport and #DRL-1418181, PI Stamper) and the Institute of Education Sciences (Training Grant R305B110003 to Liu). Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSF or IES.

## 7. REFERENCES

[1] Aleven, V., Sewall, J., McLaren, B.M., and Koedinger, K.R. (2006). Rapid authoring of intelligent tutors for real-world and experimental use. In *Proceedings of the 6th ICALT*. IEEE, Los Alamitos, CA, pp. 847-851.

[2] Baker, R.S., Corbett, A.T., Roll, I. and Koedinger, K.R. (2008). Developing a generalizable detector of when students game the system. *UMUAI*, 18(3), pp. 287-314.

[3] Barnes, T. (2005). The q-matrix method: Mining student response data for knowledge. In *Proceedings of the AAAI-EDM Workshop,* pp. 978–980.

[4] Blikstein, P. (2013). Multimodal learning analytics. In *Proceedings of the 3rd International Conf. on LAK*. ACM, New York, NY, pp. 102-106.

[5] Cen, H., Koedinger, K.R., and Junker, B. (2006). Learning Factors Analysis: A general method for cognitive model evaluation and improvement. In *Proceedings of the 8th International Conference on ITS,* pp. 164-175. Berlin: Springer-Verlag.

[6] Davenport, J., Rafferty, A., Timms, M., Yaron, D., and Karabinos, M. (2012). ChemVLab+: Evaluating a virtual lab tutor for high school chemistry. In *International Conf. of the Learning Sciences*.

[7] D'Mello, S.K., and Calvo, R. (2011). Significant Accomplishments, New Challenges, and New Perspectives. In R. A. Calvo and S. D'Mello (Eds), *New Perspectives on Affect and Learning Technologies*. New York: Springer, pp. 255-272.

[8] Graesser, A.C., Conley, M., and Olney, A. (2012). Intelligent tutoring systems. In K.R. Harris, S. Graham, and T. Urdan (Eds.), *APA Educational Psychology Handbook: Vol. 3. Applications to Learning and Teaching*. Washington, DC: American Psychological Association, pp. 451-473.

[9] Gonzalez-Brenes, J.P., and Mostow, J. (2012). Dynamic Cognitive Tracing: Towards Unified Discovery of Student and Cognitive Models. In *Proceedings of the 5th International Conf. on EDM*.

[10] Koedinger, K.R., Baker, R.S., Cunningham, K., Skogsholm, A., Leber, B., and Stamper J. (2010). A Data Repository for the EDM community: The PSLC DataShop. In Romero C, Ventura S, Pechenizkiy M, Baker RSJd (Eds.), *Handbook of Educational Data Mining*. Boca Raton, FL: CRC Press.

[11] Koedinger, K.R., Stamper, J.C., McLaughlin, E.A., and Nixon, T. (2013). Using data-driven discovery of better cognitive models to improve student learning. In *Proceedings of the 16th International Conf. on AIED.*

[12] Lan, A.S., Studer, C., Waters, A.E., and Baraniuk, R.G. (2014). Sparse Factor Analysis for Learning and Content Analytics. *Journal of Machine Learning Research,* 15, pp. 1959-2008.

[13] Lindsey RV, Khajah M, Mozer MC. (2014). Automatic discovery of cognitive skills to improve the prediction of student learning. In *Advances in Neural Information Processing Systems, 27,* pp. 1386-1394.

[14] Picard, R., Vyzas, E., and Healey, J. (2001). Toward machine emotional intelligence: Analysis of affective physiological state. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 23(10), 1175-1191.

[15] Porayska-Pomsta, K, Mavrikis, M, D'Mello, S.K., Conati, C., and Baker, R. (2013). Knowledge elicitation methods for affect modelling in education. *IJAIED*, 22, 107–140.

[16] Schneider, B., and Blikstein, P. (2014). Unraveling Students' Interaction Around a Tangible Interface Using Gesture Recognition. In *Proceedings of the 7th Annual International Conf. on EDM.*

[17] Stamper, J. and Koedinger, K.R. (2011). Human-machine student model discovery and improvement using DataShop. In *Proceedings of the 15th International Conf. on AIED*.

[18] Worsley, M. (2012). Multimodal learning analytics: enabling the future of learning through multimodal data analysis and interfaces. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction*, pp. 353-356.