

Empowering Instructors in Learning Management Systems: Interactive Heat Map Analytics Dashboard

Michael Ginda Naren Suri Andreas Bueckle Katy Börner

Cyberinfrastructure for Network Science Center

Indiana University

Wells Library, 1320 E. Tenth Street

Bloomington, IN 47405, USA

(812) 855-3259

mjinda@indiana.edu

ABSTRACT

Learning analytics visualizations can empower teachers to keep track of student engagement and performance of hundreds of students. This paper starts with a brief review of learning analytics dashboard design and the user needs of instructors. A heuristic assessment of Canvas LMS course analytics dashboards identifies limitations of current visualizations and suggests the design of a multi-level heat map of student engagement and performance. The heat map is implemented using student trace data generated by 1000 students taking the 2015 information visualization course at IU. Data selection and preprocessing workflows and dashboard visualization design are detailed. We present results of a user study involving university instructors and discuss implications for design improvements. The paper concludes with a discussion of opportunities for learning analytics dashboard development and assessment.

CCS Concepts

Applied computing → **E-learning** • **Human-centered computing** → **Visualization design and evaluation methods** • **Human-centered computing** → **Visual analytics** • **Human-centered computing** → **Heat maps**

Keywords

Visual analytics dashboards, Canvas Learning management system, e-learning

1. INTRODUCTION

One of the more challenging aspects of running and managing courses, be it residential courses with elements that take place online or a massive open online courses (MOOCs), is the ability to support students' efforts to achieve their educational goals. This is particularly true for courses that enroll students with vastly different goals and needs as common in MOOCs. Instructors need learning analytics tools and visualizations that help them provide effective support so students stay engaged, achieve learning objectives, and manage course performance and administrative tasks required for reporting by their institution [15].

With over a decade of development, learning analytics dashboards that monitor the activity and performance of students are a standard feature of any learning management systems (LMS) and other virtual education environments (e.g., MOOC platforms, intelligent tutoring systems). However, there are many opportunities to improve dashboard visualizations with the goal of improving instructor's ability to use learning analytics dashboards and visualizations effectively.

Different instructors prefer different learning analysis and visualization tools. Some pick basic statistics and tables without any visualization. Others like to use more advanced visualizations that support correlation analysis, course flow visualizations, automated clustering and classifying students [20]. The variability in instructor preference is caused by both the challenge in accessing and using student data from complex LMS data models and from the lack of abilities to using data mining or statistical analysis technique with unfamiliar tools [13, 19].

Modern LMS now offer tools to evaluate student activity and performance through dashboards that provide visual representations of the current and historical state of learners and courses data to support flexible decision making by instructors and learners [7, 4]. A review of LA dashboards showed that designs have focused on supporting instructors or collaboration between instructors and students, representing the grade performance and efforts of learners using traces of student activity [21].

Instructors can track student progress over time using a variety of data sources to understand learning outcomes, social interactions in discussion forums, and study behaviors [14-16]. Instructors in smaller online courses place more emphasis on tracking concept mastery and individual performance when compared to MOOC instructors; however, instructors generally showed preference for visualizations that were informative and easy to use [2].

Prior work argues that LA dashboard design should encompass specific goals that seek to trigger user behaviors and actions [15]. Evaluations of LA dashboards predominantly focus on the usefulness and usability of dashboards, with few dashboard evaluations looking at the efficiencies of visualization designs in support of user task completion or the effectiveness for developing or improving instructor soft skills, e.g., improving teaching or student learning and performance [21].

Dashboard designers have begun to produce heuristic guidelines that support the design of learner focused LA dashboards that promote linking student engagement to their learning outcomes [3]. While the focus of the guidelines is on student outcomes, adaptations for instructors and course designers is appropriate with a focus on supporting or developing reflective and interpretive skills of instructors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

LAK'17, March 13–17, 2017, Vancouver, BC, Canada.

Copyright 2017 ACM 1-

DOI: <http://dx.doi.org/XXXXXXXXXX>

The variability in instructor preference is caused by both the challenge in accessing and using student data from complex LMS data models and the lack of data mining and visualization tools that are easy to use [13,19].

2. ASSESSING CANVAS' COURSE ANALYTICS DASHBOARD

A heuristic assessment of the Canvas course analytics dashboard was performed to understand its utility for analyzing the engagement and performance of students within a course and to identify potential interventions for poorly performing and engaging students. Instructure's Canvas provides real-time exploratory learning analytics visualizations for instructors and students through the learner, assignment, and course analytics dashboards [12]. The instructor dashboard uses an in-memory design to load limited, separate data sets and to render modular statistical visualizations. The visualizations address specific administrative questions and allow users to toggle between the information visualizations and underlying, aggregated data [1]. The dashboard has two views: a course overview and an individual student view; see Canvas user documentation [10, 11]. The design of the two analytic do not interact, but use similar modular visualizations and basic data aggregation techniques and with inconsistent visual representation and reference systems. The disunity of the dashboard components designs and functionality limit usability and make instructor reflection and exploration of student activity difficult. Descriptions of both course analytic dashboard views follow, along with an assessment of the current designs' limitations.

2.1 Course Overview

The course overview dashboard uses three modular visualizations to help an instructor answer administrative questions: the relative activity during course, the level of assignment and assessment submissions, and the overall grades. The visualizations aggregated data for all students enrolled in the course using various temporal organizations; and there is inconsistent use of temporal scales across the visualizations. Figure 1 shows the overview course analytics dashboard for IVMOOC.

The "Activity by Date" visualization is a bar graph showing the temporal distribution of student actions over the length of the course, specifically page views and count of activities (student discussion posts, homework and quiz submissions). Bar color indicates the presence or lack of student discussion activity; however, the legends indicating this distinction are only shown both conditions are present. A user must hover over a bar to reveal the periodicity of temporal bins and underlying data values.

The "Submissions" visualization uses a stacked categorical area graph to show the proportion of on-time, late, or missing submissions. The bars represent submissions, which use a temporal arrangement based on the assigned date for an assignment. The bars display submission status proportionally, and use a green-yellow-red warning color system. By hovering over a bar shows details of the underlying quantities of submission tardiness but does not link to the submitted assignments or exams.

The "Grades" visualization uses bar and whisker charts to represent grade distribution for assignments (homework assignments, quizzes and exams) using their point values rather than percentiles; submissions are displayed in order of their due dates. The visualization allows comparison of assignment grade distributions and displays a median point values for a submission; however, a percentile scaling may be of more utility.



Figure 1. Canvas Course Analytics overview dashboard with insets: A) Activity by Time, B) Submissions, C) Grades, and D) Individual student activities data table [11].

A list of students enrolled in the course shows individual course statistics, including: page views of course materials and participation in course activities lack definitions for what qualifies as a counted behavior underlying a value. Submission variables use tardiness to define the aggregated based on student assignment metadata. The table also provides a student's current score for the course. Individual student analytics views are accessed by selecting the name of a student in the table.

2.2 Individual Student View

Within Canvas, each student in the course has an individual student analytics dashboard. The individual student dashboards mirror the course overview dashboard in a number of ways. Individual student dashboard replicates all of the administrative tasks of the course analytics dashboard, and many of the visualizations and reference systems. The individual dashboard uses the same modular visualization design supported by separate in-memory data files. Figure 2 shows a view of an individual student's course analytics dashboard.

The "Activity by Date" copies the overview dashboard visualization using an individual student's data aggregated by weekly aggregated page view and participation data. A "Conversations" visualization represent communications between instructors and students on separate aligned timelines, with each message marked by a color-coded conversation balloon symbol. The timeline follows the same temporal reference system used in the "Activity by Date" module.

The "Submissions" visualization in the individual student dashboard uses a new visual reference system, with each assignment represented as separate timelines that align with the temporal reference system of the "Activity by Date" visualization. Each assignment line and symbol indicate when a student begins and submits an assignment. Symbol and categorical color indicate if an assignment was submitted on time, late, is missing, or in the future. This visualization allows an instructor to see the overall progress of a student and identify periods of inactivity quickly.

The "Grade" visualization in the individual student dashboard uses a similar bar and whisker plot; however, the plots are also used to

indicates a student's relative submission scores and overall performance (e.g., good, fair, poor). The symbols use the same categorical symbol and color reference system as used in the individual student "Submissions" visualizations.

2.3 Assessment

Currently, course analytics dashboards implement an overview and drill down for detail framework. Individual visualizations display set statistics, which are shown when a user hovers over a graphical element. There is no way to access student artifacts, like submissions, directly from the dashboard. The separate data analysis and aggregation processes are duplicated between the course and individual student dashboard views leading to a strong connection between the visualizations across dashboards rather than internal to each separate dashboard.

While instructors define grade performance measurements when they design and implement a course in an LMS, participation and course activity measurements are not well defined. The display of student performance and engagement measures make comparison across visualizations and between engagement and performance data difficult. The use of distinct data scales makes cross-visualization comparisons difficult.

Efficient comparisons of individual students are limited to the use of the Canvas' internal gradebook and the course overview dashboard's list of student's individual activity data. Visualized comparison of students is cumbersome in the current dashboard design. Data aggregation of student activity is cumulative and qualitatively different from the temporal data aggregation used in the "Activity by Date" visualizations. The cumulative aggregation makes the daily or weekly tracking and comparison of student activity and performance over time difficult. The location of the list of student activity makes navigation between the course overview and individual student dashboards cumbersome.

The Canvas course analytics dashboards do not use administrative groups of students (course sections, project groups) to present data over time. For larger courses with multiple instructors and sections, there is no easy way to determine how sections are performing and engaging in course activities to help instructors manage their resources and explore course data to identify problems and productive solutions.

A redesigned Canvas course analytics dashboard should:

- support instructor reflection, inquiries, sense-making, and course administration;
- represent data via easy to use interfaces;
- be flexible for use across various course designs, and customizable to accommodate various user preferences and different engagement and performance measures;
- present an abstracted overview of the data first, but support more detailed exploration;
- use consistent data aggregations with defined temporal periods or categorical system;
- support additional aggregations and visual representations that represent administrative sections, groups, or clusters, and allow comparisons over time;
- define statistical measurements and data mining and normalization techniques to users;
- use a consistent visual and symbolic representation system with well-defined legends;
- provide access to student artifacts for review;
- provide access to underlying data for secondary analysis



Figure 2. Canvas Course Analytics overview dashboard with insets: A) Activity by Time, B) Communications, C) Submissions, and D) Grades [10].

3. MULTI-LEVEL HEAT MAP DASHBOARD DESIGN

The improved dashboard design proposed in this paper builds on prior lines of work. Of particular relevance are visualization that support comparison across multiple data dimensions and are flexible and extensible to alternative ordinal data arrangements within one consistent user interface.

Cross mapping is one technique to visualize relationships among the entities in collections of journal articles temporally to show linkages between authors and links to past knowledge [18]. Entity groups are determined via agglomerative clustering techniques that organize data records by their similarity, and provide visual order to complex data. In a learning analytics context, cross maps may support monitor and comparison over a variety of metrics, within and across courses; or with individual or groups of students and instructors. While appropriate for smaller courses, larger courses increase visual complexity.

Mazza and Dimitrova [14-16] offer a variety of multi-dimensional scatter plots and heat maps that use data from discussion boards, instructor identified course concepts, and student access and behavioral statistics that let instructors efficiently explore data and create mental models of course activities to determine students in need of support and portions of the course that need improvement.

Halatchliyski and colleagues [8] analyzed the main paths of knowledge evolution and contributor roles using a swim lane visualization. The display helps track the roles individuals play across multiple courses in a learning environment.

3.1 The Dashboard

The proposed dashboard takes advantage of the visual efficiency of Mazza & Dimitrova's heat map designs and modifies them to represent student activity and performance data for courses with multiple sections or large enrollments through an interactive, multi-level heat map. The dashboard (see Figure 3, data was shuffled to ensure anonymity) consists of three components: the top-level heat



Figure 3. Screenshot shows the multi-level heat map of student engagement and performance data. On top is the aggregated view with weekly engagement and submission grades for all four course sections. Selecting a cell brings up engagement or score data for individual students in the lower-level heat map. Legend and description are in lower right.

map displays student activity and performance data aggregated by course sections; the lower-level is an individual student heat map that displays activity and performance aggregated for each week of the course; and the legend provides information on the visual encoding used across both levels of the dashboard. The two heat maps interact with each other; by selecting a cell in the top heat map causes the lower-level heat map to update with either weekly student activity or grade performance data that corresponds. The visualization uses a ranking normalization method, and weighted indicators to allow comparison across generalized student activity and performance behaviors.

3.1.1 Top-level heat map view

The top-level heat map displays student data is aggregated section (rows) and week (column). Columns may represent student weekly averages of student engagement (magenta to blue color scale) or performance on major graded submissions (red to green color scale) throughout the course. The dashboard allows instructors to access and compare student activity and grade performance of each section. It helps identify sections that are (not) engaging in the course activities or submitting and performing well on assignments.

3.1.2 Lower-level heat map views

The lower level of the visualization has two heat maps viewed by interacting with the top-level heat map cells. The first view of the lower level inset is generated by selecting a cell representing average student engagement in the top-level heat map. The inset shows a break out of the weekly activity for individual student across engagement statistics. Engagement is broken out into indicators, page views, discussions, quizzes. An overall engagement score and current course grade is the final column in the lower inset and is a weighted rank of all engagement factors.

To show the second view of the lower-level inset of major assignment and overall grades, a user must select a cell that presents an assignment or exam grade in the top-level heat map. The inset shows a simplified view of the individual student grades for

selected assignments, a current score for the course, and an average engagement score for the student for the student in the course.

3.1.3 Spatial layout

The spatial layout of the heat map visualization allows a viewer to see and interact with both the top and lower levels of the heat map simultaneously.

The lower level inset only uses a portion of the screen space to allow for the display of a legend that explains the color scales used for the heat map. The legend has a brief description of the visualization, and leaves room for future extensions that provide task specific statistical visualization insets and/or details on data underlying the engagement indicators when one selects a cell for an individual student.

Both heat maps display groups and individual students along the y-axis. For the top-level heat map, columns use a chronological order along the x-axis. For lower-level insets, columns use either an ad hoc or a chronological order for columns.

3.1.4 Legends and color scales

Using two color schemes helps to differentiate between types and scales of data used for student engagement and grade performance indicators. Cells within the heat maps representing student activity and engagement indicator rankings use a diverging color scale to indicate a clear distinction between active, moderate, and inactive students in the course. The colors transition on a linear scale assigned to a given percentile rank value, between magenta to white to blue; magenta represents low or no activity and blue indicates very active students. Cells representing student grade performance or major submissions also use a diverging color scale to indicate grade performance. Colors use a traditional scheme that transition on a linear scale between red to yellow to green; red represents low or no scores and green indicates very high scores for a student. The dashboard represents these two scales as gradient bars in the legend with associated percentile and course letter grading scales.

3.2 Information Visualization MOOC

The multi-level heat map dashboard visualizes student activity and performance data for the Information Visualization MOOC (IVMOOC), taught each spring since 2013. Students from more than 100 countries may freely take the course with graduate students also take the course for credits towards their degree at Indiana University. The course provides an overview about the state of the art in information visualization, and covers data temporal, geospatial, topical, and network analysis algorithms and visualization techniques that enable extraction of patterns and trends, and discussions of systems that drive research and development. For the first half of the course, theoretical lectures and hands-on tutorials ground students work to explore temporal, geospatial, topical, and network analysis and visualization techniques. The second half of the course asks students to collaborate in teams on information visualization projects and collaboration with real-world clients. In spring 2015, four sections—one free IVMOOC and three bearing IU credits—were taught yet all students share the same resources (lecture and tutorial videos) and activities (homework assignments, self-assessment quizzes, exams, discussion forums, and client projects).

3.3 Data and Processing Workflow

The IVMOOC course uses the Canvas learning management system. The Canvas Data product collects course activity data. Instructure compiles and delivers this data to customers through Amazon Web Services via the Redshift data warehouse. The Redshift data warehouse uses a star database schema to differentiate between descriptive dimensions and qualitative factors of various aspects of a course. Canvas Data was accessed and retrieved in December 2015 using SQL and database administration tool connected to the Redshift data warehouse. Using the “course_id”, student’s “canvas_id” was collected from the “user_dim” table, and student sections, then using both sets of identifiers, student sections were collected from “enrollment_dim” tables. Next, web logs collect traces of student activity across factors in the “requests” table; discussion data collects in “discussion_entry_dim” and “discussion_topic_dim” tables; and grades collect in “submission_dim” table [8]. Retrieved data were stored in a local PostgreSQL database of learning analytics data related to the IVMOOC course, with enrollment table identifiers acting as the primary key to query associated course weblogs, discussion posts, and gradebook data.

3.3.1 Calculating Indicators

3.3.1.1 Selecting enrollments

Information on student enrollments is required to examine data quality and to analyze and visualize student activity. A list of all enrolled student was created by selecting those accounts with the enrollment type “student” and deleting all known test enrollments (used by instructors to test IVMOOC functionality). Only student accounts that showed activity—at least one web log—during the first two weeks of the course were kept all others discarded. The result are exactly 1000 student enrollments.

3.3.1.2 Grade Performance Data

Next, student grades were collected from the submissions data in the Redshift database, including grades for exams, assignments. Individual scores for these assignments were preserved for visualization and to calculate students’ cumulative grades. Student enrollment section and student identifiers were used to join the data with weekly student activity indicators into a single data set.

3.3.1.3 Student Activity Data

Student activity was aggregated for individual students’ data to create set of indicators ranking student activity for first nine weeks of the course and the weeks preceding the start of the course to show if students were active prior to the start of the course. Indicators of student engagement include discussions, page views, and quizzes. These engagement variables were compiled using SQL and Python scripts using data collected from Canvas Redshift data warehouse as follows:

Weekly discussion indicators are calculated using a weighted value model based on variables calculated used data found discussion posts and replies tables. The discussions indicator based on the count of discussions topics started by a student, the count of replies made by a student, and a calculation of the average length of their posts for each week.

Weekly page view counts were calculated for each student by filtering the requests web logs to count only records with the value “GET” in the “HTTP METHODS” field. The filter implemented approximates the individual page views automatically calculated by Canvas; however, a precise definition for tabulating this variable from Requests web logs was unavailable.

The quiz indicator uses two statistics based in analysis of the Requests web logs to calculate the number of quiz attempts and number of questions and attempted. Counting the number of records with a “QUIZ_ID” field and a “GET” value for the “HTTP_METHOD” field provides the number of quiz attempts; the number of question attempted counts changes made per quiz by counting web log by counting “PUT” and “POST” in the “HTTP METHODS” field. The number of correct answers from the quiz grades submissions table.

3.3.2 Data Analysis and Normalization

A descriptive analysis provided insight into the quality and composition of the student trace data, and revealed distributions of variables calculated from the Canvas Redshift data warehouse.

A temporal analysis and plot of the request log data using weekly counts of page views showed that the web controller did not properly calculate for the first two weeks of the course and for two days during the first half of the course. Specifically, data from the first two weeks of the course lacked the “HTTP_METHOD” field, preventing accurate counts of page views for this period. In addition, unfiltered page views include noise from system pings that do not indicate student generated activity. A generalized linear regression model allowed for approximation of page view counts for the first two weeks using data from the third through fifth week of the course. The model compared students’ weekly page view counts calculated from the raw weblogs and a filtered version of the weblogs that used the “HTTP Method” field.

Normalization of data using percentile ranking creates bins for activity measures as ranks while minimizing the effects of outliers in the data, and makes comparisons of indicators more efficient. For the creation of the linear regression model, page views values were normalized by assigning a percentile ranking value between zero and one. Normalization was later applied to other student activity indicators (composed of one or more related and weighted variables). The calculation of percentile ranks uses cumulative temporal periodicity, meaning all data from preceding weeks of the course are included in the creation of new rankings. Cumulative time windows weigh historic student activity in the course while accounting for fluctuations in activity and the loss of active students.

3.3.3 Weighted Value Model

Weighted value models are used to create both statistical indicators of related statistical variables (e.g., discussion variables, or quiz variables) and to calculate an overall engagement indicator composed of individual student activity indicators. The overall student engagement I_{overall} is defined as the sum of weighted (N) student engagement activities (I_x) per week:

$$I_{\text{overall}} = (1/N)*I_1 + (1/N)*I_2 + \dots + (1/N)*I_N$$

Modifications to the general weight model should be a feature for instructors to modify, to ensure variables underlying indicators do align with task needs and pedagogical goals. For example, the discussion indicator uses a modified weighted value model, where discussions ($I_{\text{discussions}}$) accounts for a student's total number of topics started (D_{started}), total number of replies (D_{replies}), and the average length of the text of posts (L_{avg}) in a given week using the following formula:

$$I_{\text{discussions}} = (1/2)*D_{\text{started}} + (1/4)*D_{\text{replies}} + (1/4)*L_{\text{avg}}$$

The modifications provide extra weight to students who start new discussions. The resulting metric is the normalized as a percentile rank using a cumulative time slice to assign values to weekly activity.

3.3.4 Data Compilation

A single data file supports the multi-level student engagement and performance heat map. The JSON xml file comprises weekly student engagement and score data for each student. It is organized temporally with a hierarchical data structure for sets of related fields. For week of the course combines the student activity indicators with students' current overall engagement indicator. Additional sets of columns are created for major assignments or exams in the course, and follow the week a submission was due. These additional sets of columns represent each major graded assignment submissions the course, the current score for the course, and the current overall engagement indicator rank for each student.

3.4 Deployment

The dashboard displays in web browsers using D3 and AngularJS, and CSS. The two levels of the heat map are visualized simultaneously as separate insets within a screen and use the same data source, which allows for dynamic updates of the lower-level when a user interacts with the top-level visualization.

4. USER STUDY

This study examines the readability of the multi-level heat map for the 2015 Information Visualization course by semi-experts. Test subjects had to be current or former instructors, and must have used a learning management system (LMS) to be eligible for the study. Six testers participated in the study. This following section details study setup, participants, and data analysis results.

4.1 Participants and Setting

The study asked participants to complete a pre-questionnaire to capture information on basic demographics and any previous experience with data visualizations or tools to create them.

Next, participants viewed and interacted with the dashboard visualization in a web browser on a computer and given a task sheet with instructions and about a dozen quantitative and qualitative questions. The task sheet had two sections: "Course Section View" contained questions and prompts for the aggregated view while "Individual Student View" was concerned with the view of individuals in the sections.

There were tasks with a precise answer, such as "Which section has the highest engagement score?", and prompts, such as "What is the key insight you get from exploring this section of the visualization?" This allowed us to a) get feedback for the further development process of the tool and b) made the participants interact and play with the tool so they could learn more about its functionality.

4.2 Data Analysis and Results

The majority of participants (n=6) were or are associate instructors or teachers or PhD students and in the age group of 21-30, with one exception (one participant was records assistant for the university). All subjects but two were affiliated with the School of Informatics and Computing (SOIC). The majority (n=5) were male. All but one were English native speakers, with one participant speaking Chinese/Cantonese as their first language. No participants have had prior training in visualizations, but two used a wide array of visualization software before (e.g., Jupyter Notebook).

When asked to explore the visualization and answer questions, there were certain tasks that all the participants got right: most importantly, the fact that section Z637-44781 had the highest overall scores, and the IVMOOC section had the lowest, a fact established through four questions about the aggregated view.

Results from multiple testers indicate that information retrieved from the top-level aggregated view was more accurate than that retrieved from the individual student view. For example, estimations about the number of students in section Z637-32593 vary from 30 to 50. Asked about the percentage of active students, the answers go from 5% to 96%. Five out of six participants, however, determined that student #9 had the lowest active page view percentage. The individual student view visualization seemingly makes it hard to estimate aggregate numbers.

In addition, we received a lot of feedback on the overall design of the tool. From the first question that asked the students to write down what the visualization shows, all participants indicated the title of the tool, which seems to be descriptive and informative. Testers drew a number of key insights from the top-level view, included were general insights ("lower student engagement in a course correlates with lower test scores"; "Higher engagement reflects on the scores of students in a positive way"), as well as more specific ones ("IVMOOC had lowest engagement and lowest scores were from section Z637-33781. More engagement seems to indicate higher scores on upcoming exams."; "33781 consistently outperforms while IVMOOC consistently underperforms"). When asked about what they liked about the visualization design, four out of six indicated that they like the color scheme. However, one tester criticized that the dark blue/purple gradient was harder to grasp compared to the red/green scheme. When asked about what the metrics in the individual view mean, the definitions given by the participants varied a bit (e.g., "[...] Engagement = average of the two" vs. "[...] engagement: actually using some part of the web page"). Two participants criticized the lack of a sorting function for the columns in the individual view. Finally, asked about how to improve the visualization, there were three themes to the answers: (1) Work on style ("Alternative color schemes available for those with different types of color-blindness"; "In dark background color, use white font"). (2) Add a sorting function for the columns in the individual view ("Also sorting columns and better yet, allowing selection of multiple sections from the aggregate view"). (3) Add mean scores ("Add an average engagement tab to the top graph").

5. DISCUSSION

The design of the multi-level heat map visualization sought to improve upon current course analytics dashboards available in the Canvas learning management system across a number of criteria, see listing in Section 2.3. The current design allows instructors to examine and compare course data across course sections and between individual students with consistent data aggregation methods, symbolic representation, and access to detailed student engagement and performance data. The current interface is useable by instructors. However, the results of the user study (see Section 5) suggest diverse improvements.

The current design does not allow for access to underlying data, e.g., exam or assignment results; does not let instructors define new statistical measurements or normalization techniques; and it does not yet provide access to underlying data for secondary analysis. However, these components may be added in planned extensions of the dashboard, or as secondary dashboards that answer other questions, such as using discussion data to extract and analyze social networks mined and perform sentiment and topic analysis.

Future iterations of the multi-level heat map will integrate data from other data science courses at Indiana University. This work will involve automation of indicator calculations and data aggregations that align with current learning analytics data standards, such as Caliper. Adding other course data will support the generalization and validation of the design across pedagogical styles and different course structures, different numbers of students enrolled in the course, and course level.

We plan to support testing preferences for instructors on customization options for layouts, indicator selection and calculations. For example, course dashboards may support alternative students aggregations based on grades and engagement indicator ranges, students' demographic groups (e.g., country, age group, gender, education) or administrative data (e.g., major, GRE scores) provided by the university. Aggregations might also represent machine learning outputs generated using student demographic and or learning outcomes.

In sum, the design of dashboards for instructors provides many challenges for designers of LMS. Dashboards must be easy to use by users with limited data analysis and visualization skills or they will not succeed in improving online learning experiences. Close collaboration of users, data mining and visualization experts, learning scientists, and commercial providers (e.g., developers of canvas or Unizin's SnapShot) are essential for designing dashboards that support smart decision making by instructors.

6. ACKNOWLEDGMENTS

Our team would like to thank Sara Bouchard and Adam Simpson for their expert work on the design and implementation of the visualizations; the Indiana University Information Technology Services Canvas team, including Matt Gunkel, Randall Embry, Tan Nguyen, Justin Zemlyak, Lynn Ward, Erik Scull, and David Goodrum, for providing access to Canvas data; and Linda Shepard of the Bloomington Assessment and Research group for providing additional student data and IRB support. Technical information and support on Canvas web analytics and Redshift data store was provided by Matt Smith and Chris Ward of Instructure. This project was funded by the Google Education and University Relations Fund of Tides Foundation and by an Indiana University Bloomington Learning Analytics Fellows Program Award.

REFERENCES

- [1] Abdelfattah, M. 2013. A comparison of several performance dashboards architectures. *Intelligent Information Management*, 5, 35-4. DOI= doi:10.4236/iim.2013.52005
- [2] Ali, L., Asadi, M., Gašević, D., Jovanović, J., and Hatala, M. 2013. Factors influencing beliefs for adoption of a learning analytics tool: An empirical study. *Computers and Education*, 62, 130-148. DOI= <http://doi.org/10.1016/j.compedu.2012.10.023>
- [3] Charleer, S., Klerkx, J., Duval, E., De Laet, T., & Verbert, K. 2016. Creating effective learning analytics dashboards: Lessons learnt. In K. Verbert, M. Sharples, & T. Klobočar Eds. *Adaptive and Adaptable Learning: 11th European Conference on Technology Enhanced Learning*, (Lyon, France, September 13-16, 2016), EC-TEL, 42-56. DOI= http://dx.doi.org/10.1007/978-3-319-45153-4_4
- [4] Emmons, S., Light, R. and Börner, K. 2016 [In Press]. MOOC visual analytics: Empowering teachers, students, researchers, and developers of massively open online courses. *Journal of the Association for Information Science and Technology*.
- [5] Duval, E. 2011. Attention please!: Learning analytics for visualization and recommendation. In *Proceedings of the 1st International Conference on Learning Analytics and Knowledge*, (New York, NY, USA), LAK '11, ACM, New York, NY. 9-17. DOI= <https://doi.org/10.1145/2090116.2090118>
- [6] Dyckhoff, A. L., Zielke, D., Bültmann, M., Chatti, M. A., and Schroeder, U. 2012. Design and implementation of a learning analytics toolkit for teachers. *Educational Technology & Society*, 15 (3), 58-76.
- [7] Few S. 2006. *Information dashboard design: The effective visual communication of data*. Sebastopol, CA: O'Reilly Media.
- [8] Halatchliyski, I., Hecking, T., Goehnert, T., and Hoppe, H. U. 2014. Analyzing the main paths of knowledge evolution and contributor roles in an open learning community. *Journal of Learning Analytics*, 1 (2), 72-93. Retrieved from <http://epress.lib.uts.edu.au/journals/index.php/JLA/article/view/3534>
- [9] Instructure. 2016. *Canvas data portal. Schema documentation, [version 1.13.0]*. Technical documentation. Retrieved from <https://portal.inshosteddata.com/docs>
- [10] Instructure. 2016. *How do I view analytics for a student in a course?* User documentation. Retrieved from <https://guides.instructure.com/m/4152/1/66791-how-do-i-view-analytics-for-a-student-in-a-course>
- [11] Instructure. 2016. *How do I view Course Analytics?* User documentation. Retrieved from <https://guides.instructure.com/m/4152/1/66790-how-do-i-view-course-analytics>
- [12] Instructure. 2016. *What are course analytics?* User documentation. Retrieved from <https://guides.instructure.com/m/4152/1/66789-what-are-course-analytics>
- [13] Krüger, A., Merceron, A., & Wolf, B. 2010. A data model to ease analysis and mining of educational data. In M. Pechenizkiy et al. eds, *Proceedings of the 3rd International Conference on Educational Data Mining 2010* (Eindhoven,

The Netherlands), International EDM Society, Pittsburg, PA, 131–140.

- [14] Mazza, R., and Dimitrova, V. 2007. CourseVis: A graphical student monitoring tool for supporting instructors in web-based distance courses. *International Journal of Human-Computer Studies*, 65(2), 125–139. DOI=<http://doi.org/10.1016/j.ijhcs.2006.08.008>
- [15] Mazza, R., and Dimitrova, V. 2004. Visualising student tracking data to support instructors in web-based distance education. In *WWW Alt. '01* (New York, NY, May17-22, 2004), ACM, New York, NY, 154. DOI=<http://doi.org/10.1145/1013367.1013393>
- [16] Mazza, R., and Dimitrova, V. 2003. Informing the design of a course data visualisator: an empirical study. In *5th International Conference on New Educational Environments* (ICNEE 2003), 215–220.
- [17] Morris, S. A., and Boyack, K. W. 2005. Visualizing 60 years of anthrax research. *Proceedings of the 10th international conference of the international society for scientometrics and informetrics*. (Stockholm, Sweden, July 24-28), Karolinska University Press, Stockholm, SE 45–55.
- [18] Morris, S. A., and Yen, G. G. 2004. Crossmaps: Visualization of overlapping relationships in collections of journal papers. *Proceedings of the National Academy of Sciences*, 101(suppl 1), 5291–5296. DOI=<http://doi.org/10.1073/pnas.0307604100>
- [19] Pedraza-Perez, R., Romero, C., & Ventura, S. 2011. A Java desktop tool for mining Moodle data. In M. Pechenizkiy et al. (Eds.), *Proceedings of the 3rd Conference on Educational Data Mining* (Eindhoven, The Netherlands, 2011), International EDM Society, Pittsburg, PA, 319–320.
- [20] Stephens-Martinez, K., Hearst, M. A., and Fox, A. 2014. Monitoring MOOCs: Which information sources do instructors value? In *Proceedings of the First ACM Conference on Learning @ Scale Conference*, (New York, NY, USA), ACM, New York, NY. 79–88. DOI=<http://doi.org/10.1145/2556325.2566246>
- [21] Verbert, K., Duval, E., Klerkx, J., Govaerts, S., & Santos, J. L. (2013). Learning analytics dashboard applications. *American Behavioral Scientist*, 2764213479363. DOI=<https://doi.org/10.1177/0002764213479363>