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

Michael Ginda, Naren Suri,
Andreas Bueckle, and Katy Börner
Cyberinfrastructure for Network Science Center, School of Informatics, Computing, and Engineering, Indiana University Bloomington

CNS

Learning analytics visualizations can empower teachers to keep track of student engagement and performance of hundreds of students. This poster provides 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 Indiana University. 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 poster concludes with a discussion of opportunities for learning analytics dashboard development and assessment.

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 [5].

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

Prior work argues that LA dashboard design should encompass specific goals that seek to trigger user behaviors and actions [5]. 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.

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. 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.

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 [4,6].

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.

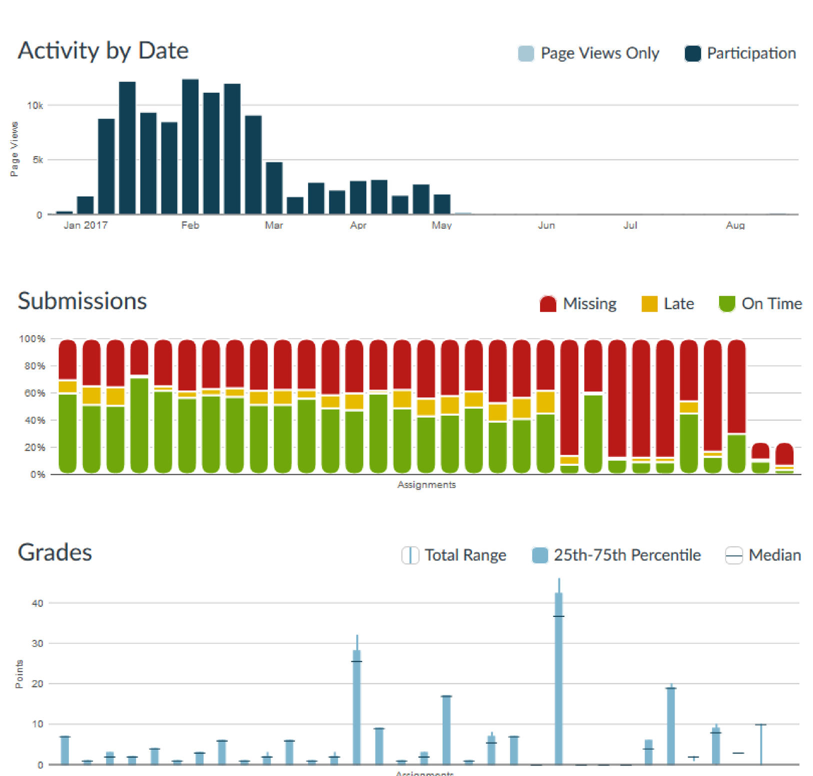

Figure 1. Canvas Course Analytics overview dashboard with insets: A) Activity by Time, B) Submissions, C) Grades, and D) Individual student activities data table [3].

2.1 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.

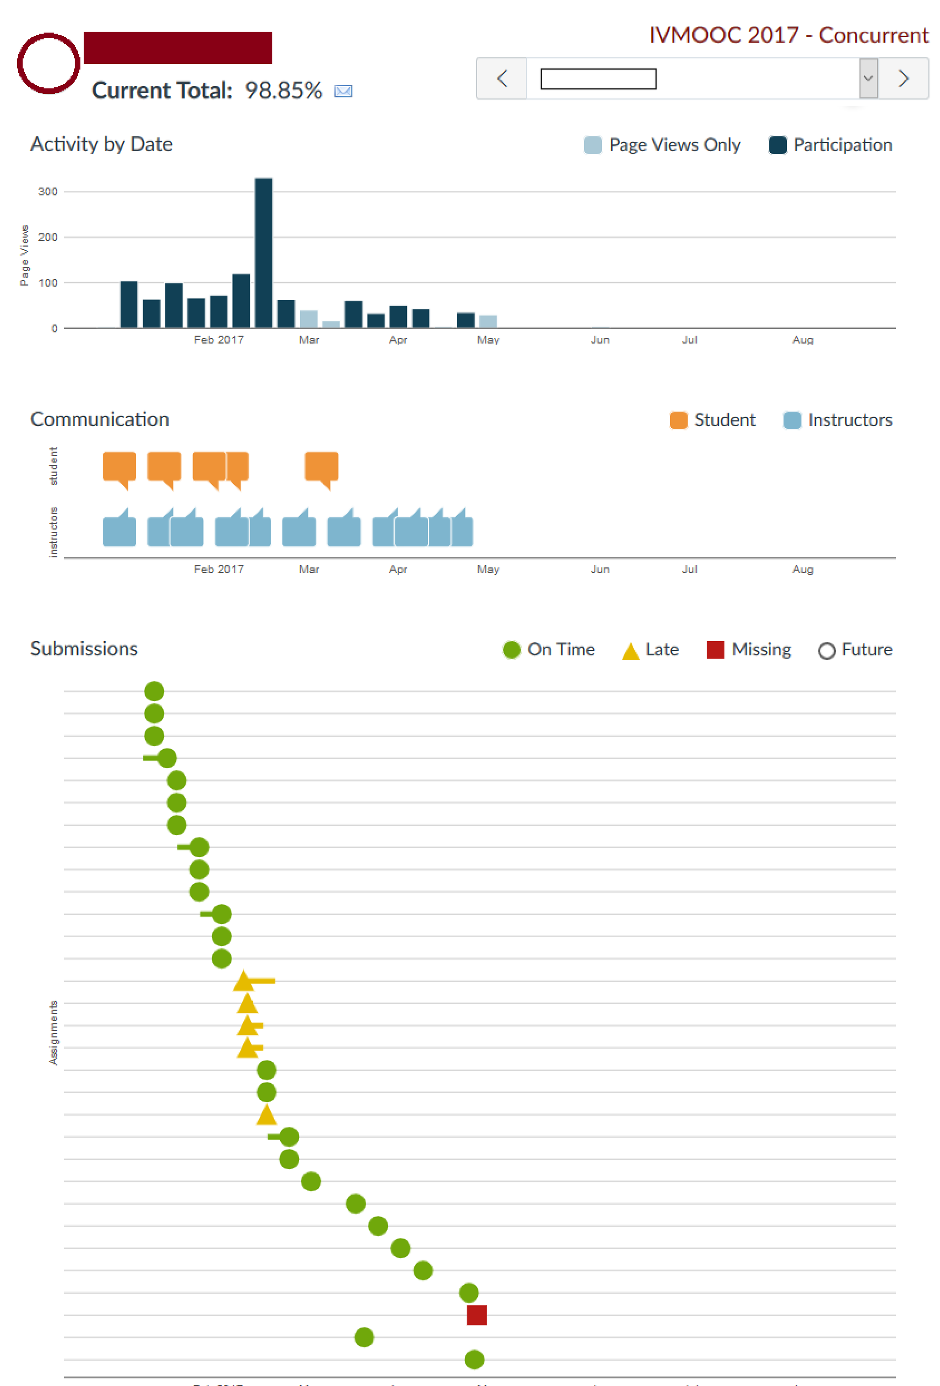

Figure 2. Canvas Course Analytics overview dashboard with insets: A) Activity by Time, B) Communications, C) Submissions, and D) Grades [3].

2.2 Assessment

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

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.

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 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 data 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.

3.2 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).

3.4 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.

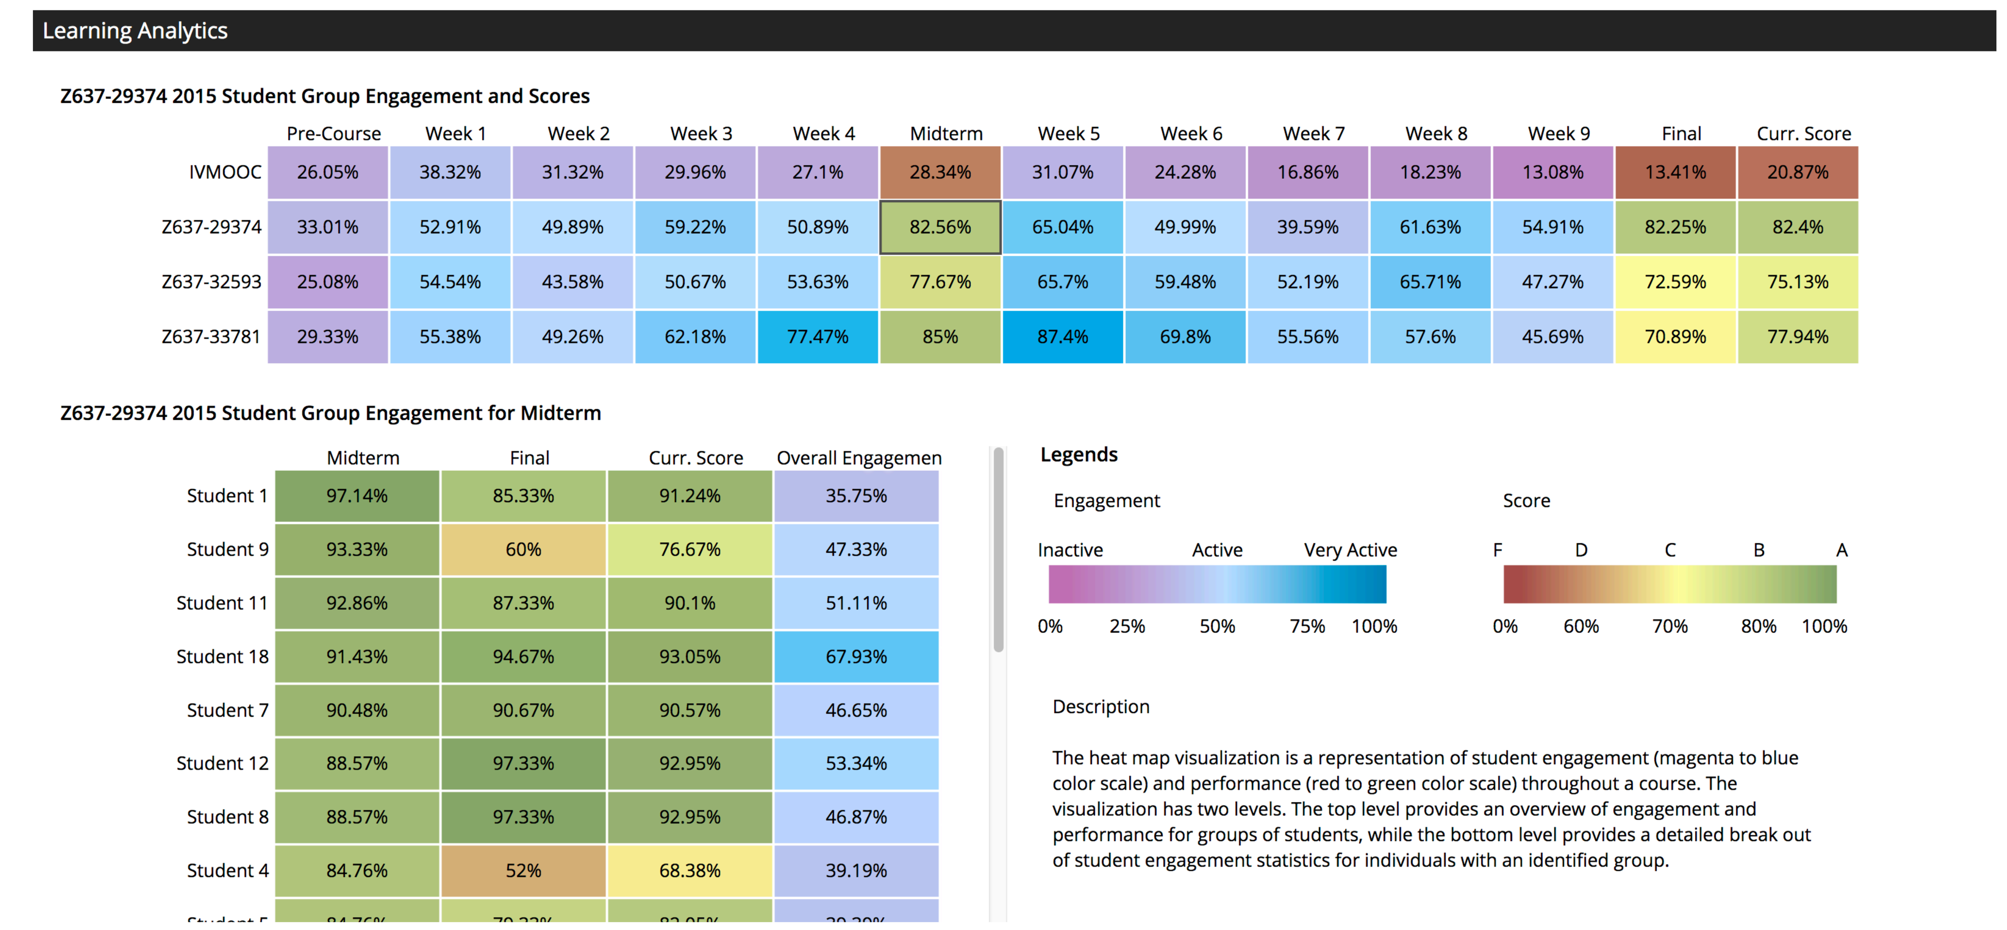

Learning Analytics

Z637-29374 2015 Student Group Engagement and Scores

	Pre-Course	Week 1	Week 2	Week 3	Week 4	Midterm	Week 5	Week 6	Week 7	Week 8	Week 9	Final	Curr. Score
IVMOOC	26.05%	38.32%	31.32%	29.96%	27.1%	28.34%	31.07%	24.28%	16.86%	18.23%	13.08%	13.41%	20.87%
Z637-29374	33.01%	52.91%	49.89%	59.22%	50.89%	82.56%	65.04%	49.99%	39.59%	61.63%	54.91%	82.25%	82.4%
Z637-32593	25.08%	54.54%	43.58%	50.67%	53.63%	77.67%	65.7%	59.48%	52.19%	65.71%	47.27%	72.59%	75.13%
Z637-33781	29.33%	55.38%	49.26%	62.18%	77.47%	85%	87.4%	69.8%	55.56%	57.6%	45.69%	70.89%	77.94%

Z637-29374 2015 Student Group Engagement for Midterm

	Midterm	Final	Curr. Score	Overall Engagemen
Student 1	97.14%	85.33%	91.24%	35.75%
Student 9	93.33%	60%	76.67%	47.33%
Student 11	92.86%	87.33%	90.1%	51.11%
Student 18	91.43%	94.67%	93.05%	67.93%
Student 7	90.48%	90.67%	90.57%	46.65%
Student 12	88.57%	97.33%	92.95%	53.34%
Student 8	88.57%	97.33%	92.95%	46.87%
Student 4	84.76%	52%	68.38%	39.19%

Legends

Engagement

Inactive Active Very Active

0% 25% 50% 75% 100%

Score

F D C B A

0% 60% 70% 80% 100%

Description

The heat map visualization is a representation of student engagement (magenta to blue color scale) and performance (red to green color scale) throughout a course. The visualization has two levels. The top level provides an overview of engagement and performance for groups of students, while the bottom level provides a detailed break out of student engagement statistics for individuals with an identified group.

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 flour 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.

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 4) suggest diverse improvements.

References

[1] 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. Klobučar 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

[2] Instructure. 2016. *How do I view analytics for a student in a course?* User documentation. Retrieved from https://guides.instructure.com/m/4152/l/66791-how-do-i-view-analytics-for-a-student-in-a-course

[3] Instructure. 2016. *How do I view Course Analytics?* User documentation. Retrieved from https://guides.instructure.com/m/4152/l/66790-how-do-i-view-course-analytics

[4] 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.

[5] 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

[6] 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

[7] 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