Leveraging Peer-Learning Environments for Visualizing Engagement, Achievements and Competencies at Scale

Hassan Khosravi
University of Queensland
h.khosravi@uq.edu.au

Kendra Cooper
Independent Scholar
kendra.m.cooper@gmail.com

ABSTRACT

Learning dashboards present a large amount of data using visualisations that support exploration of learning activities by different stakeholders. Students can use the visualisations for self-reflection and comparison of themselves with their peers, and teachers can use the visualisations to gain insight on class-level and individual-level knowledge gaps. Most of the current learning dashboards rely on logs as their primary data source, which greatly reflects students' engagement level, but they are not a reliable source for capturing the achievements or competencies of learners. In this paper, a new learning dashboard for Visualising Engagement, Achievements, and Competencies at Scale (VEACaS) is presented. This learning dashboard harnesses data generated through students' interactions with a Peer-Learning Environment, called PeerWise, and provides visualisations that demonstrate the level of engagement, achievements, and competencies of the students. VEACaS is analysed using both synthetic and real data sets. Results are promising, indicating that VEACaS can provide meaningful visualisations to both students and teachers in different learning environments.

Keywords

Learning Dashboards, Visual Analytics, Peer-Learning Environments, Competencies

1. INTRODUCTION

As class sizes grow and classroom environments evolve (traditional, e-learning, flipped, MOOCs), educators continue to face significant challenges in effectively engaging, assessing, and providing high-quality, personalised feedback to students, while working within limited budgets for scarce, highly skilled staff. Recently, peer-learning has emerged as a scalable approach to actively engage, assess, and provide timely feedback. For example, PeerWise [8] is a free web-based system in which students can both create multiple-choice questions, and answer, rate, and discuss questions created by their peers. Empowering students with environments like these offers well-recognised, significant benefits that include enhancing student involvement in cognitively demanding tasks such as identifying missing knowledge, diagnosing misconceptions, and providing feedback in their own words. From a practical perspective adopting Peer-Learning Environments in large classes generates extensive, rich data on students based on their interactions with the system; these data have the potential to yield deeper insights to further help students.

To help make sense of data collected through students' interactions with educational tools and technologies, an active research area in the learning analytics community focusing on visual analytics and learning dashboards has emerged [22, 12, 15]. The work on visual analytics and learning dashboards present a large amount of data by visualisations that supports exploration of learning activities by students and teachers. Students can use the visualisations for self-reflection and comparison of themselves with their peers and teachers can use the displays to gain insight on class-level and individual-level knowledge gaps. Once the gaps are identified, teachers can update their course content and improve the quality of assessments and the provided feedback to address those knowledge gaps.

Here, preliminary results on VEACaS are presented. This dashboard harnesses data generated through students' interactions with PeerWise, and provides visualisations that effectively communicate the level of engagement, achievements, and competencies of the students. VEACaS is organised into three main components: *Input Data* that stores students' interactions with the Peer-Learning Environment; *Data Integration* that computes students' engagement, approximates their achievements, and infers their competencies; and *Visualisations* that visualise students' engagement, achievements and competencies.

The behavior of VEACaS is first validated and examined under different circumstances using synthetic data sets, in which the underlying knowledge gaps of the students are predefined. We then apply VEACaS to analyse real data from an on-campus introductory course in C programming for engineering students at The University of British Columbia.

The remainder of this paper is organised as follows: a review of the related work and literature is presented in Section 2 and an overview of the VEACaS is presented in Section 3. The experimental results are reported in Section 4, and finally, conclusion and future work are presented in Section 5.

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

2. RELATED WORK

The field of visualisation has been utilised broadly to allow users to employ a variety of visual displays to explore and interpret their data [14]. With the increase in the use of educational technologies and the advancements in the areas of learning analytics and educational data mining, a new field, commonly known as "Learning Dashboards" has emerged to help make sense of data sets in learning and education [18]. This field is rapidly evolving and growing; [25] reviewed nine major learning dashboards that have been introduced through academic journals and international conferences, and [15] published a recent comprehensive survey that compared various aspects of learning dashboards from 55 papers. Based on their review, they introduce the following definition for a learning dashboard: "A learning dashboard is a single display that aggregates different indicators about learner(s), learning process(es) and/or learning context(s) into one or multiple visualisations."

Data Sources. The most common form of student data analysed is details of operations or activities that are executed by the learners, which are generally referred to as logs. Numerous recent studies used logs from a variety of environments including Learning Management Systems [27], MOOCs [19], lectures [20], and social networks [26] to analyse learner's behavior. Furthermore, [15] reported that 85% of the reviewed papers mentioned logs as their primary source of data in their learning dashboard. Although logs represent the engagement level of learners accurately, they are not a reliable source for capturing the achievements or competencies of learners [1]. In contrast, data generated through the use of Peer-Learning Environments is often rich and contains detailed information about the students' competencies.

Visualisations. Typically, in learning dashboards, a variety of traditional visualisations such as bar charts [7], pie charts [28], histograms [21], box plots, scatter plots [25] and timelines [13] have been used. Open learner models often use different visualisation forms such as skill meters, concept maps, and tree structures [11, 3]. To the best of our knowledge there is no research that has confirmed which visualisations are preferred for learning dashboards, and in fact, there is little consensus about which are most effective. In VEACaS, box and whisker plots are used to communicate the results. They provide more statistical information than traditional bar charts allowing the students to better compare themselves to the cohort. In addition, they can be easier to comprehend compared to some of the non-traditional visualisations, as stakeholders are familiar with them.

Target Users. Leaning dashboards are designed for four main types of target users: students [7, 6, 23], teachers [2, 4, 29], administrators [17, 16], and researchers [24]. The majority of the current dashboards support only a single type of target user [15]. VEACaS supports both students and teachers.

3. OVERVIEW OF VEACaS

At a high level, VEACaS applies a suite of established approaches to harness data available in Peer-Learning Environments and provides visualisations demonstrating each users' engagement, achievements, and competencies. VEACaS is organised into three main modules: Input Data; Data Integration; and Visualisations. Figure 1 provides an overview.

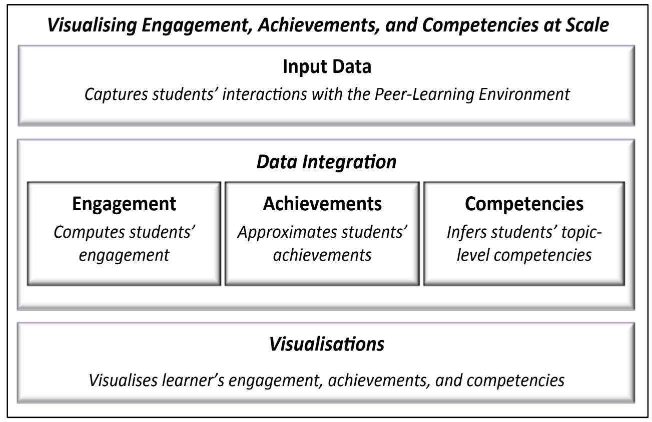

Figure 1: Overview of VEACaS for Peer-Learning Environments.

In what follows, Section 3.1 provides more information about the input data; Section 3.2 discusses how these data are then used to compute engagement, approximate achievements, and infer competencies; finally, Section 3.3 discusses the visualisations.

3.1 Input Data

Let N denote the number of users that are registered on the Peer-Learning Environment, M for the number of multiple-choice questions that have been contributed to the environment, and L as the number of distinct topics that have been used to tag the questions. Data input from PeerWise is passed on to VEACaS using three tables/matrices:

Participations, P: A summary of user's participations including a record of questions they have authored, answered, evaluated, and commented on as well as information on badges they have received are represented in a table P, where P_u captures the participation summary for user u.

Answers, A: The correctness of the answers provided by the users is represented in a matrix $A_{N\times M}$. If a user u answers a question i correctly then the matrix entry is set at $a_{ui} = 1$, $a_{ui} = 0$ indicates an incorrect answer, and $a_{ui} = Null$ indicates that question i has not been attempted by user u.

Tags, T: Each question can have 0 to L topics assigned (i.e. tagged) to it. The information on topics assigned to each question is represented in a matrix $T_{M\times L}$, in which $t_{ij} = 0$ indicates that question i is not tagged with topic j and $t_{ij} = \frac{1}{g}$ indicates that question i is tagged with $1 \leq g \leq L$ associated topics, including j.

3.2 Data Integration

The input data are used to compute engagement, approximate achievements, and infer competencies as described in the following subsections.

3.2.1 Engagement

Table P is used to compute students' engagement into the following four categories:

Answers Submitted show the number of questions that have been answered by each student.

Questions Authored show the number of questions that have been authored by each student.

Comments Written show the number of comments that have been written by each student. The comments are associated with questions.

Badges show the number of badges that have been awarded to each student. There are a total of 25 distinct badges that may be earned.

3.2.2 Achievements

Table P is used to approximate students' Achievements using two scores, Answer Score and Reputation Score:

Answer Score is based on learner's own behavior and correlates positively with the number of questions they answer correctly and negatively with the number of questions answered incorrectly. The punishment received for answering questions incorrectly is %20 of the reward provided for answering questions correctly. This is to encourage students to answer more questions while trying to discourage students from just randomly guessing.

Reputation Score rewards learners for making contributions that are valued by their peers, so unlike the answer score, it is computed based on the actions of their peers instead of their own. The reputation score is based on agreement of a reviewer with other peer reviewers who have reviewed the same material and is designed with the goal of promoting early participation and fair evaluation of other questions [5].

Both the Answer Score and Reputation Score are computed by PeerWise and are readily available for visualisation.

3.2.3 Competencies

This module uses the input data T and A to produce a student-topic learning profile $LP_{N\times L}$, in which each vector (lp_u) infers a user's competencies across all of the topics associated with the course. A positive value in the vector, i.e. $lp_{uj} > 0$ indicates that the user u has demonstrated some knowledge on topic j, a negative value indicates a knowledge gap on that topic, and 0 represents a neutral state, where the positive and negative scores have balanced each other out for that particular topic. The learning profile is computed in two steps:

1. Matrix $A_{N\times M}$ stores information about the correctness of the answers provided by the users and matrix $T_{M\times L}$ stores information about the tags associated with each question. Multiplying the two ($A_{N\times M} \times T_{M\times L}$) allows for an understanding to be gained about topic-level competencies and knowledge gaps in the system *per se*. The value stored in cell $[u, j]$ of the resulting matrix depends on the number and weight of questions tagged with topic j that have been attempted by each user u. This means that the values in this matrix require normalization.
2. Normalization is achieved using a user-topic count matrix $C_{N\times L}$, in which c_{uj} represents the weighted sum of questions attempted by user u that have been tagged with topic j. This matrix can be computed using $C_{N\times L} = S_{N\times M} \times T_{M\times L}$, in which s_{ui} is 1 if question i is attempted by user u and 0 otherwise.

Putting the two steps together, the learning profile is computed using the following formula:

$$LP_{N\times L} = \frac{G_{N\times M} \times T_{M\times L}}{C_{N\times L}} \qquad (1)$$

3.3 Visualisations

A total of 16 (four for engagement, two for achievements, and ten for competencies) indicators are visualised in VEACaS. All of the indicators are visualised using the same type of graph, box and whisker plots. The reason for using the same type of graph is so users spend less time trying to comprehend different graph types. Box and whisker plots provide statistical information on several main features that allow students to better compare themselves to the cohort and at the same time are easier to comprehend compared to some of the non-traditional visualisations. The end of the whiskers can represent several possible alternative values. In the representation used in VEACaS, the end of the whiskers show the lowest datum within 1.5 IQR of the quartile. This representation (1) allows the plots to better show the trends of the class without being affected by outliers and (2) places a higher emphasis on displaying potential outliers allowing for further examination.

Engagement level of the learners is measured using number of answers submitted, questions authored, comments written, and badges earned as described in Section 3.2.1. Achievements of the learners are measured Answer Score and Reputation Score as described in Section 3.2.2. Finally, competencies are measured based on the learning profile LP as described in Section 3.2.3.

4. EXPERIMENTAL EVALUATION

One of the goals of VEACaS is to provide visualisations that help students recognize their competencies and knowledge gaps, which in turn can help them better plan their studies. It also provides teachers with an overview of class-level competencies and knowledge gaps that allows them to adapt their course content accordingly. In what follows, we will first use synthetic data sets to examine and analyze the inferred competencies under different circumstances, in which the underlying competencies and knowledge gaps of the students are predefined. We will then utilise a historical data set for the exploration of learning activities in an introductory programming course.

4.1 Generating Synthetic Data Sets

The experiments discussed in this section make use of synthetic data sets generated using the following sequence of steps. First, pre-defined class-level knowledge gaps over five topics are created. The class-level knowledge gaps are constructed by sampling from a Dirichlet distribution, where α defines the sparsity of the distribution; a smaller value of α creates a sparser distribution over knowledge gaps, simulating a class with a large gap over one topic. The class-level knowledge gaps are stored in a vector g, where g_l represents the class-level knowledge gap of topic l. The sum of the values in g equal to 1. Second, a set of users with predefined knowledge gaps over the five topics are created. The user knowledge gaps are constructed by sampling from a Dirichlet distribution parameterised by a vector $g\times\gamma$, where $1 \leq \gamma \leq 10$ determines the diversity in students' competencies relative to g; larger values of γ leads to less diversity in students' competencies. Third, a set of questions with a pre-defined topic, level of difficulty and discrimination are generated. The topic associated with a question is sampled from a discrete uniform distribution, and the level of difficulty and discrimination are both sampled from a normal distribution.

The probability of a user u answering a question i correctly is computed using a two parameter logistic Latent Trait Models (IRT) model from classical Item Response Theory [10], as recommended by [9]:

$$\frac{1}{1+e^{-a_i(\theta_s - b_i)}} \tag{2}$$

where θ_s represents user's lack of knowledge gap (competency) in the topic of question i, b_i is the difficulty level and a_i is the discrimination level of question i.

In all generated data sets 400 users, 1100 questions, and 22000 answers are sampled, which approximates the numbers from the historical data set that is used for exploration in Section 4.4.

4.2 Impact of Varying Pre-Defined Class-level Competencies

Figure 2 illustrates the impact of α, which defines the sparsity of pre-defined class-level competencies. Use of small values of α generate a sparse distribution over knowledge gaps, which simulates classes with a large knowledge gap on a single topic. Figure 2a visualises competencies for a data set where $\alpha = 0.01$, in which the class-level knowledge gap is on topic 2. This example resembles courses that have one or more topics that the majority of the students find challenging and often perform poorly on. Updating the course material with particular attention to the topic(s) with the class-level knowledge gap is an effective way to improve such courses.

Increasing α generates less sparse distributions, which leads to the simulation of courses with less extreme class-level pre-defined knowledge gaps. Figure 2b visualises competencies for a data set where $\alpha = 1$, in which the class-level knowledge gap is still on topic 2; however, the gap is not as large as the previous case. This example resembles courses that have topics that many of the students find challenging and often perform poorly on. Updating the course material on the topic(s) that are the class-level knowledge gap as well providing personalised feedback for students on their strengths and weaknesses are effective ways of improving such courses.

Use of large values of α generates a discrete uniform distribution over knowledge gaps, which simulates classes with knowledge gap over all of the topics. Figure 2c visualises competencies for a data set where $\alpha = 100$, in which topic three seems more challenging than the other topics; however class-level knowledge gaps are almost uniformly distributed. This example resembles courses in which students' knowledge gaps are very different. Providing personalised feedback for students on their strengths and weaknesses are effective ways of improving such courses.

4.3 Impact of Varying Diversity of Students' Competencies

Figure 3 illustrates the impact of γ, which defines the diversity in students competencies relative to the class-level pre-defined knowledge gaps, g. The underlying class-level knowledge gap distribution does not have an impact on the analysis of γ. In this example g is arbitrarily chosen to be the following $\{0.14, 0.85, 0.05, 0.05, 0.01\}$.

Use of large values of γ generates synthetic users with knowledge gaps similar to the class-level pre-defined knowledge gaps, which simulates classes with little diversity among students. Figure 3a visualises competencies for a data set where $\gamma = 100$, in which students' knowledge gaps are similar to g; topic 2 is the main gap. This example resembles small courses that are only open to students in a particular program. Assuming similarity among the students means that there is less benefit in personalization in such courses.

Decreasing γ, generates synthetic users with a larger deviation from the class-level knowledge gaps, which leads to simulation of courses with more diverse students in terms of pre-defined knowledge gaps. Figure 3b visualises competencies for a data set where $\gamma = 10$, in which students' knowledge gaps are still similar to g; however, there is more diversity in students' knowledge gaps and competencies as indicated by the larger interquartile ranges in the box plots. This example resembles larger courses that are open to a wider range of students. The larger difference among students means that there is more benefit in personalization in such courses.

Use of small values of γ, generate synthetic users which are only mildly affected by g and are mostly sampled from a discrete uniform distribution over knowledge gaps, which simulates classes with large diversity among students. Figure 3c visualises competencies for a data set where $\gamma = 1$, in which the knowledge gap in the class is still topic 1; however, there is a significant difference in students' knowledge gaps and competencies as indicated by the even larger interquartile ranges in the box plots. This example resembles Massive Open Online Courses (MOOCs), which are open to wide range of self learners. The large difference among students means that there are huge benefits to provide personalization in such courses.

4.4 Historical Data Set

Having examined the behavior of the model with different class-level competencies and varying levels of diversity using synthetic data sets, we now turn to analysing the use of VEACaS with a historical PeerWise data set. The data set is from a required, introductory course in C programming for engineering students, offered at The University of British Columbia during 2014. To encourage participation, students received grades for their use of the PeerWise environment: (i) They were required to author at least three questions and to answer at least 45 questions (worth 1.5% of final mark), and (ii) a grade was calculated from the "Answer Score" (AS) and " Reputation Score" (RS), which are computed by the PeerWise system as described in Section 3.2.2, using the following formula: $\frac{min(AS,RS)\times 1.5}{500}$, (worth 1.5% of final mark). In total, 377 students authored 1111 questions and answered 21432 questions, assigning 1700 tags to the questions, based on 10 topics.

Figure 4 illustrates VEACaS presented to a randomly chosen student; We use the nickname Alex to refer to this student.

The general class-level patterns reveal that most students are well engaged with PeerWise in this course. The median number of questions answered by students is 57, which is significantly higher than the requirement for receiving full-grades, which is to answer 45 questions. The median number

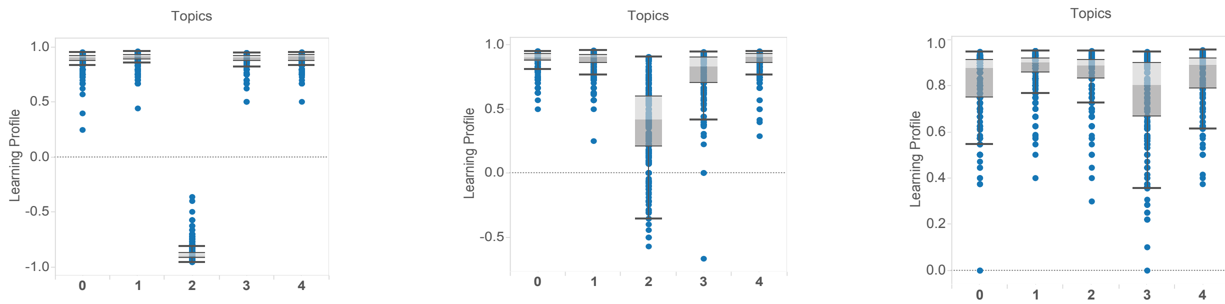

(a) α is equal to 0.01 (b) α is equal to 1 (c) α is equal to 100

Figure 2: Impact of varying pre-defined class-level competencies. As α is increased the diversity among class-level knowledge gaps is increases.

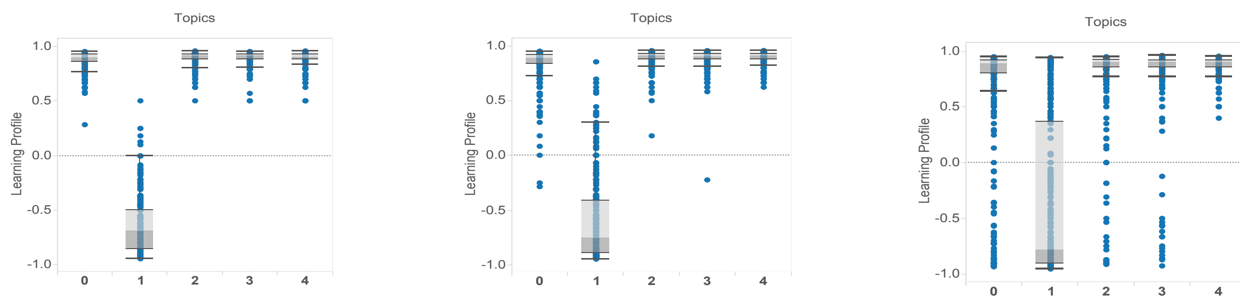

(a) γ is equal to 100 (b) γ is equal to 10 (c) γ is equal to 1

Figure 3: Impact of varying diversity of students' competencies. As γ is decrease the diversity among competencies of students increases.

of authored questions is equal to the requirement for obtaining full-grade, which is three questions. There were no direct requirements for writing comments, and therefore, most students have not engaged in writing comments. The Reputation Scores of the cohort has a median of 1695 and a lower quartile of 1011, which is much higher than the minimum requirement for obtaining full-grade, which is 500. The Answer Scores of the cohort has a median of 370 and a lower quartile of 240, which is much closer to the minimum requirement for obtaining full-grade, which is 300.

The class-level competencies reveal that there are students with and without knowledge gaps on all of the topics (i.e, the computed competencies in all of the topics range from a negative number to a positive number). Programming comprehension, conditionals, DAQ system, fileIO, functions, and programming syntax topics have positive medians indicating that the majority of the students in the class had a positive value in their learning profile for these topics. Loops and fundamentals have negative medians, indicating that the majority of the students have a negative value in their learning profile for these topics. Finally, arrays and introduction have medians close to zero indicating that roughly half the students had positive values and the other half had negative values for these topics in their learning profile. An interesting observation is that the whiskers, quartiles, and the median in the programming comprehension topic are all 0.4. This consistency can be explained by the limited number of questions, 15 out of 1111, that are on the programming comprehension.

Looking closely at Alex's interactions with the PeerWise Environment, we can see that overall, Alex's engagement is low compared to the cohort. Alex has answered and authored fewer questions than 75% of the students in the class; Alex has answered 35 questions while the lower quartile of the class is 42 and has authored one question while the lower quartile of the class is two. Furthermore, Alex has not written any comments, where the median of the class is one; he has received six distinct badges, which is equal to the lower quartile of the class.

In terms of achievements, which are approximated using the Reputation Score and the Answer Score that are computed by PeerWise, Alex's achievements are lower than 75% of the students in the class. Alex's Reputation Score is 750, where the lower quartile of the class is 1011 and Answer Score is 218, where the lower quartile of the class is 240.

Despite having poor engagement and achievements, Alex has demonstrated a high level of competency on some of the topics in the course. For example, Alex has performed extremely well (i.e., has a score, which is higher or equal to the upper quartile of the class) on questions on programming comprehension, DAQ systems, functions, and programming syntax. Alex has done relative well (i.e., has a score higher or equal to the median but lower than the upper quartile) on questions on conditionals, FileI/O and Fundamentals. Based on the provided answers, Alex is performing poorly (i.e., has a score lower than the lower quartile) on questions on Arrays.

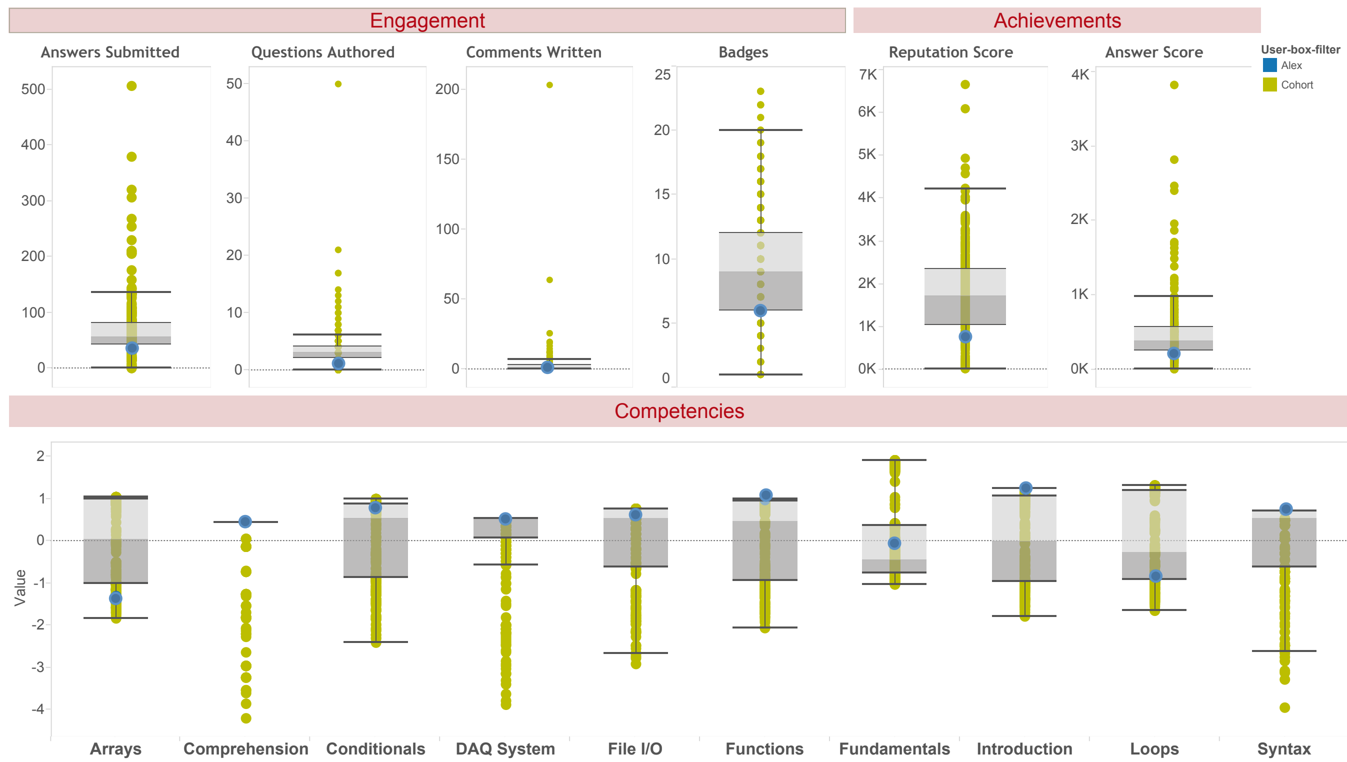

Figure 4: VEACaS from the view of a randomly selected student nicknamed Alex. The learning dashboard visualises Alex's engagement, achievements, and competencies.

5. CONCLUSION AND FUTURE WORK

Supporting exploration of learning activities by different stakeholders at scale is a challenging task. In this paper VEACaS, a learning dashboard for Visualising Engagement, Achievements, and Competencies at Scale is presented. Unlike most learning dashboards that primarily rely on logs, which capture students' engagements, VEACaS visualises students' competencies, achievements, and engagement.

VEACaS consists of three main modules: Input Data; Data Integration; and Visualisations. The Data Integration module uses the output of the Input Data module to compute engagement, approximate achievements, and infer competencies of the students. Competencies are measured by multiplying matrix A, which stores the correctness of the answers provided by the students, and matrix T, which captures the topics associated with each question. The output is a student-topic learning profile in which each vector infers a user's competencies across all of the topics associated with the course.

Experimental validation of the dashboard used both synthetic data sets and a historical data set. The synthetic data sets were used to examine the inferred competencies in the dashboard under diverse circumstances, by varying parameters of the data generation template. Different parameter settings resembled different learning environments from classes that are open only to students within a particular program to MOOCs, which are open to wide range of self-learners. The historical data set was used for exploration of learning activities in an introductory programming course.

There are several interesting directions to pursue in future work: (1) The formulation of the learning profile can be modified, so that the significance of the gap may be affected by the difficulty level of the questions. The difficulty level of a question can either be approximated by computing the ratio of correct answers to total answers or by explicitly allowing users to rate the difficulty level of the question, which is supported by PeerWise.

(2) The current dashboard visualises a user's knowledge gaps but doesn't have the capacity to assist the user in overcoming the gaps. The dashboard would benefit from a recommendation engine that can recommend questions that are most likely to help the user overcome their knowledge gap. Ideally, the recommender system would be able to prioritise questions in which the user is likely to have an interest, as this will help to maintain their engagement with the system.

6. REFERENCES

[1] Á. F. Agudo-Peregrina, Á. Hernández-García, and S. Iglesias-Pradas. Predicting academic performance with learning analytics in virtual learning environments: A comparative study of three interaction classifications. In *Computers in Education (SIIE), 2012 International Symposium on*, pages 1–6. IEEE, 2012.

[2] S. Aguilar, S. Lonn, and S. D. Teasley. Perceptions and use of an early warning system during a higher education transition program. In *Proceedings of the fourth international conference on learning analytics and knowledge*, pages 113–117. ACM, 2014.

[3] S. Bull, B. Ginon, C. Boscolo, and M. Johnson. Introduction of learning visualisations and metacognitive support in a persuadable open learner model. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge*, pages 30–39. ACM, 2016.

[4] S. Charleer, J. L. Santos, J. Klerkx, and E. Duval. Improving teacher awareness through activity, badge and content visualizations. In *International Conference on Web-Based Learning*, pages 143–152. Springer, 2014.

[5] K. Cho and C. D. Schunn. Scaffolded writing and rewriting in the discipline: A web-based reciprocal peer review system. *Computers & Education*, 48(3):409–426, 2007.

[6] L. Corrin and P. de Barba. Exploring studentsâĂŹ interpretation of feedback delivered through learning analytics dashboards. In *Proceedings of the ascilite 2014 conference*, pages 629–633, 2014.

[7] S. Dawson, L. Macfadyen, E. F. Risko, T. Foulsham, and A. Kingstone. Using technology to encourage self-directed learning: The collaborative lecture annotation system. 2012.

[8] P. Denny, J. Hamer, A. Luxton-Reilly, and H. Purchase. Peerwise: students sharing their multiple choice questions. In *Proceedings of the fourth international workshop on computing education research*, pages 51–58. ACM, 2008.

[9] M. C. Desmarais and I. Pelczer. On the faithfulness of simulated student performance data. In *Educational Data Mining 2010*, 2010.

[10] F. Drasgow and C. L. Hulin. Item response theory. 1990.

[11] D. Duan, A. Mitrovic, and N. Churcher. Evaluating the effectiveness of multiple open student models in eer-tutor. 2010.

[12] E. Duval. Attention please!: learning analytics for visualization and recommendation. In *Proceedings of the 1st International Conference on Learning Analytics and Knowledge*, pages 9–17. ACM, 2011.

[13] A. L. Dyckhoff, D. Zielke, M. Bültmann, M. A. Chatti, and U. Schroeder. Design and implementation of a learning analytics toolkit for teachers. *Educational Technology & Society*, 15(3):58–76, 2012.

[14] U. M. Fayyad, A. Wierse, and G. G. Grinstein. *Information visualization in data mining and knowledge discovery*. Morgan Kaufmann, 2002.

[15] A. Holzer, A. Voznuik, D. Gillet, M. J. Rodríguez-Triana, B. A. Schwendimann, L. P. Prieto, M. S. Boroujeni, and P. Dillenbourg. Perceiving learning at a glance: A systematic literature review of learning dashboard research. *IEEE Transactions on Learning Technologies*, 2016.

[16] S. Išljamović and S. Lalić. Academic dashboard for tracking students' efficiency. In *proceedings of the Xiv International Symposium Symorg 2014: New Business Models and Sustainable Competitiveness*, page 84. FON, 2014.

[17] E. Kapros and N. Peirce. Empowering l&d managers through customisation of inline learning analytics. In *International Conference on Learning and Collaboration Technologies*, pages 282–291. Springer, 2014.

[18] P. A. Kirschner, S. J. Buckingham-Shum, and C. S. Carr. *Visualizing argumentation: Software tools for collaborative and educational sense-making*. Springer Science & Business Media, 2012.

[19] R. F. Kizilcec, C. Piech, and E. Schneider. Deconstructing disengagement: Analyzing learner subpopulations in massive open online courses. In *Proceedings of the Third International Conference on Learning Analytics and Knowledge*, LAK '13, pages 170–179, New York, NY, USA, 2013. ACM.

[20] J. M. Kraushaar and D. C. Novak. Examining the affects of student multitasking with laptops during the lecture. *Journal of Information Systems Education*, 21(2):241, 2010.

[21] D. Leony, A. Pardo, L. de la Fuente Valentín, D. S. de Castro, and C. D. Kloos. Glass: a learning analytics visualization tool. In *Proceedings of the 2nd international conference on learning analytics and knowledge*, pages 162–163. ACM, 2012.

[22] R. Mazza and V. Dimitrova. Visualising student tracking data to support instructors in web-based distance education. In *Proceedings of the 13th International World Wide Web Conference on Alternate Track Papers &Amp; Posters*, WWW Alt. '04, pages 154–161, New York, NY, USA, 2004. ACM.

[23] K. Muldner, M. Wixon, D. Rai, W. Burleson, B. Woolf, and I. Arroyo. Exploring the impact of a learning dashboard on student affect. In *International Conference on Artificial Intelligence in Education*, pages 307–317. Springer, 2015.

[24] Z. A. Pardos and K. Kao. moocrp: An open-source analytics platform. In *Proceedings of the Second (2015) ACM conference on learning@ scale*, pages 103–110. ACM, 2015.

[25] Y. Park and I.-H. Jo. Development of the learning analytics dashboard to support students' learning performance. *J. UCS*, 21(1):110–133, 2015.

[26] E. Popescu. Providing collaborative learning support with social media in an integrated environment. *World Wide Web*, 17(2):199–212, 2014.

[27] C. Romero, S. Ventura, and E. García. Data mining in course management systems: Moodle case study and tutorial. *Computers & Education*, 51(1):368–384, 2008.

[28] A.-M. Tervakari, K. Silius, J. Koro, J. Paukkeri, and O. Pirttila. Usefulness of information visualizations based on educational data. In *Global Engineering Education Conference (EDUCON), 2014 IEEE*, pages 142–151. IEEE, 2014.

[29] A. Vozniuk, M. J. Rodríguez-Triana, A. Holzer, S. Govaerts, D. Sandoz, and D. Gillet. Contextual learning analytics apps to create awareness in blended inquiry learning. In *Information Technology Based Higher Education and Training (ITHET), 2015 International Conference on*, pages 1–5. IEEE, 2015.