Incorporating Explainable Learning Analytics to Assist Educators with Identifying Students in Need of Attention

Shiva Shabaninejad, Hassan Khosravi, Solmaz Abdi, Marta Indulska, Shazia Sadiq
The University of Queensland, {shiva}uq.edu.au
Brisbane, QLD, Australia

Abstract

Increased enrolments in higher education, and the shift to online learning that has been escalated by the recent COVID pandemic, have made it challenging for instructors to assist their students with their learning needs. Contributing to the growing literature on instructor-facing systems, this paper reports on the development of a learning analytics (LA) technique called Student Inspection Facilitator (SIF) that provides an explainable interpretation of students learning behaviour to support instructors with the identification of students in need of attention. Unlike many previous predictive systems that automatically label students, our approach provides explainable recommendations to guide data exploration while still reserving judgement about interpreting student learning to instructors. The insights derived from applying SIF in an introductory Information Systems course with 407 enrolled students suggest that SIF can be utilised independent of the context and can provide a meaningful interpretation of students' learning behaviour towards facilitating proactive support of students.

CCS Concepts

• Information systems → Data analytics; **• Applied computing → Computer-assisted instruction**; **E-learning**.

Keywords

Explainable Learning Analytics, Learning analytics dashboards, at-risk Students

ACM Reference Format:
Shiva Shabaninejad, Hassan Khosravi, Solmaz Abdi, Marta Indulska, Shazia Sadiq. 2022. Incorporating Explainable Learning Analytics to Assist Educators with Identifying Students in Need of Attention. In *Proceedings of the Ninth ACM Conference on Learning @ Scale (L@S '22), June 1–3, 2022, New York City, NY, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3491140.3528292

1 Introduction

The growing popularity of remote and online learning in higher education, together with a significant and rapid increase in online learning due to the COVID pandemic, have made it challenging for instructors to realise and attend to the academic needs of their students. The LA community has presented various approaches to address this issue. Promising approaches include the use of learning analytics dashboards (LADs) that provide visualisations to help instructors with sensemaking regarding what is happening in their course [21, 28], early warning systems that automatically identify at-risk students [3, 7], tools that support instructors with sending personalised feedback [5, 23], as well as newly developed LA methods that incorporate state of the art process mining and machine learning algorithms to infer student learning tactics and strategies [2, 16] or to recommend sub-populations of students who deviate from the rest of the class based on learning process metrics [19, 30]. Despite the advancements these solutions provide, each comes with its own challenges and limitations: Many LADs have been reported to have a low level of impact as they mostly present descriptive statistics over raw clickstream data without any grounding to learning theories or contextualisation, which makes the interpretation and decision making difficult [17]. Likewise, many early warning systems are under increasing scrutiny due to the concerns about the fairness, accountability, transparency, and ethics (FATE) of black-box predictive models that provide lists of the at-risk students without any explanation or justification [15]. On the contrary, personalised feedback tools that are based on manual filtering for personalised feedback, such as OnTask [23], and SRES [5], enable instructors to provide explainable and actionable feedback for students. However, they require a notable time commitment from instructors and rely on their complex multi-dimensional data navigation skills, which have been shown to be a challenge for instructors without a technical background [29]. There is also a range of newly developed LA methods and tools that address the issues of LADs and early warning systems by providing in-depth knowledge about students learning behaviour from temporal data [2, 16, 30] but they are context-dependent and not suitable for adoption at the scale of an institute with multiple teaching disciplines.

To mitigate some of the shortcomings of the current approaches, in this paper, we present the Student Inspection Facilitator (SIF) - a LA approach to help educators identify students in need of attention. SIF is designed to meet the following objectives: (1) support utilisation cross disciplines and modalities, (2) support identification of students in need of attention without requiring instructors to spend an extensive amount of time and effort, (3) support AI recommendations with explainability and interpretability to allow for instructor oversight. SIF delivers a holistic view of each student and prioritises for instructors' personalised support, based on five risk factors, grounded in literature, including: (1) performance level, (2) engagement level, (3) performance consistency, (4) engagement consistency, and (5) deviation from the class norm. To demonstrate how SIF could be used in practice, we explore the insights derived from the inspection of the students' learning characteristics using

L@S '22, June 1–3, 2022, New York City, NY, USA

ACM ISBN 978-1-4503-9158-0/22/06…$15.00
https://doi.org/10.1145/3491140.3528292

data from an introductory Information Systems course with an enrolment of 407 students.

2 Design Objectives

As a basis for our design, we turn to the existing literature to inform SIF objectives.

Context Independence. In order for a LA solution to make a long-term impact, its capability to be adopted at the scale of an institute is crucial [11]. One of the critical challenges of adopting LA from an institutional point of view that is identified by literature [11, 25], is where the LA technologies are designed to be compatible with a particular educational context. For instance, a set of newly developed LA techniques utilise process mining techniques on students' engagement data to derive insights about students learning strategies. For instance, [16] successfully identifies students' tactics and strategies [2] across different learning design modalities [9]. However, implementation of this approach requires contextualisation for different course designs. Similarly, [30] uses process mining to identify sub-populations of students who deviate from the rest of the class based on their learning processes. However, the validity of the results depends on the instructional design of the course and how fine-grained the data is, which makes it unsuitable for institute-wide adoption. Accordingly, to leverage these concerns, we incorporate a core objective of *supporting utilisation cross disciplines and modalities* in the design of SIF.

Effective use of instructor time. A recognised concern with a range of existing LA techniques is that they demand a lot of time from academics [12]. For instance, descriptive LADs provide visualisations and data exploration techniques for instructors' observation and interpretation of student data [26]. Although instructors are well-positioned to verify the pedagogical soundness of findings of LA systems and contextualise them to their teaching, instructor-led efforts to manually explore complex multidimensional data sets to identify students that need attention require technical expertise and is time-consuming [29]. To mitigate the concerns regarding academic workload, we incorporate a core objective of *supporting identification of students in need of attention without requiring instructors to spend an extensive amount of time and effort* forms another core objective of SIF.

Explainability and interpretability. Early LA approaches aimed to provide insights from student data by making automatic interpretation and recommendations on students or course outcome. These approaches are mainly used to predict on students outcomes based on students' logged data and demographic data [26]. Predictive LADs generally, operate as "black-box" giving the users no insight into the rationale of their choice [1]. Black-box recommendations, have been found to be sensitive and limited to real-world observational data, which are plagued with confounders and biases [8, 15, 18]. Accordingly, to minimise automatic decision making's biases, we incorporate a core objective of *supporting AI recommendations with explainability and interpretability to allow for instructor oversight* in the design of SIF.

3 Student Inspection Facilitator (SIF)

In this section, guided by our core system objectives, we introduce SIF. SIF is an inspection dashboard that presents students, color-coded and ordered by their priority to be inspected. For each student SIF provides a holistic view of the student's data, and the interpretation of their learning behaviour (i.e. the risk factor score), to facilitate instructor's support decisions. To meet our first objective - *context independence* - SIF's interpretation of students' learning behaviour is designed to be based on a flexible set of engagement and performance metrics to allow the input to be context independent (e.g., weekly engagement) or context-specific (e.g., engagement with specific content), with no dependence on a specific disciplinary context. To meet our second objective - *effective use of instructor time* - SIF prioritises students to minimise instructors' effort of searching through the whole set of students and to help with identification of students whose learning behaviour is more uncommon to be observed at an aggregated level. To avoid automatic decision making and to engage instructors in judgement about students' learning, aligned with our third objective, SIF provides *explainability and interpretability* of the students' rankings by following Explainable artificial intelligence (XAI) principles suggested by [13, 22]. SIF's interface is designed based on the principles of show-and-tell approach that provides justifications of decisions made by the AI system accompanied by visualisations of input data used to generate inferences.

Notation and definitions. Assume that a LA system has access to a data set D containing information about a set of students $S = \{s_1, \dots s_N\}$ on a set of learning behaviour metrics such as assessment scores and engagement level in different learning activities $M = m_1, \dots m_K$. Data pertaining to student s_n is represented in row n of D and data pertaining to metric k is represented in column k of D. Therefore, d_{nk} stores information on students s_n based on metric m_k. Our goal is to use a set of instructor selected metrics $M' \in M$ and data set D to infer a set of analytics outcomes $O = \{o_1, \dots o_L\}$ that are used for ranking students on their likelihood to need attention.

3.1 Methods and Algorithms

SIF' algorithm takes two parameters as input: the dataset D and the instructor selected learning behaviour metrics M' and returns a ranked student list based on a set of risk factors, represented by Σ. In our implementation, we considered the following analytics as risk factors: (1) performance compared to the rest of the class, (2) engagement compared to the rest of the class, (3) consistency of performance across the assessment scores, (4) consistency of the engagement across the engagement scores, (5) deviation from the class norms. The algorithm consists of six main segments, which are discussed below.

- Scale the selected metrics to an equal range. Our algorithm normalises the students' scores in the selected metrics M' to allow for the comparison of heterogeneous attributes. The algorithm applies a standardisation method based on z-score transformation to scale all the metrics to the same range. The normalised scores D' are next used as an input for the four the following functions.

- Compute student performance level and engagement level. This segment of the algorithm computes performance and engagement ranking scores $0 \leq performance_s \leq 1, 0 \leq engagement_s \leq 1$ for a student s, where a higher ranking score is given to students with a lower performance or engagement. It takes the dataset D' and sets a ranking score $performance_s/engagement_s$ for each student in D' based on the average of their performance scores. For instance, $performance_s = 0.05$ indicates that student s is at 5% range of the class in terms of the poorness of their performance. Each score is evaluated relative to the class's score distribution to minimise the impact of external factors that affect the class overall behaviour (e.g., the difficulty level of assessments).

We consider these two metrics because students' online engagement with learning management systems and students' study performance, have been reported as the most commonly used source of data to identify students for intervention [32]. Correspondingly, a wide range of LA techniques use these data to identify under-engaged and under-performed students, and provide evidence on the effectiveness of intervention towards those students [6, 10, 27].

- Compute student performance and engagement consistency. This segment of the algorithm computes a performance consistency ranking score $0 \leq performanceConsistency_s \leq 1$ and an engagement consistency ranking score $0 \leq engagementConsistency_s \leq 1$ based on performance and engagement metrics M' for each student, where a higher ranking score is given to students that have a lower consistency in engagement or performance. We approach this problem using the standard-deviation function to measure the differences between the scores for each student. Similar to the previous segments, each score is evaluated relative to the class's score distribution to minimise the impact of external factors on the evaluation of students' behaviour.

We consider these metrics as it has been shown that consistency in learning behaviour and engagement reflect self-regulation strategies taken by students and is an important factor in identifying at-risk students [31, 33, 33]. While many reasons can explain inconsistency in learning behaviour, such as higher workload in other courses, it is still worthwhile to monitor and inform the students about the possible negative impact on their learning [14].

- Compute student deviation from the class norms. This segment of the algorithm computes a class deviation ranking score $0 \leq classDeviation_s \leq 1$ for a student s, where a higher ranking score is given to students whose learning process deviates the most from their peers. It takes D' as input and sets a ranking score based on each student's distance from the class norms. We approach the computation of class-deviation distance as an outlier-detection problem. To measure the outlier score, our algorithm uses the X-Means clustering method to infer an appropriate number of centroids [24] and to find coherent groups in the students' set. Then, the distance of each students' score is computed based on (local outlier factor) mechanism [4] using euclidian-distance function as $\sqrt{\sum(p_i - q)^2}$ where p,q are two points in Euclidian n-space.

We consider this metric to identify students with the most different learning processes, a factor in previous LA systems that report on the most deviated subgroups of students [29, 30]. We argue that this outlier data may provide useful insights that identify misconduct cases or students in need of support.

- Rank students based on their likelihood of needing attention. Finally, we compute an overall ranking score for students based on their likelihood to require attention using the resulting scores in the five risk factors. To obtain a ranking score, we make use of stochastic dominance [20] from decision theory where one set of outcomes (set of risk factor scores for student s) can be considered superior to another set of outcomes (set of risk factor scores of student s'). The aim is to order the students in such a way that one student is less 'at-risk' than another if their distribution of outcomes is both smaller on average and less variable, while students at a higher-risk have larger scores on average and more variability. For obtaining our function, we consider the following two criteria, which are related to the first-order and second-order stochastic dominance:

(1) The overall ranking score for student s dominates the overall ranking score for student s' if each risk score of student s dominates that of scores of student s'. If $(performance_s) > (performance_{s'})$, $(engagement_s) > (engagement_{s'})$, $performanceConsistency_s > performanceConsistency_{s'}$, $engagementConsistency_s > engagementConsistency_{s'}$ and $classDeviation_s > classDeviation_{s'}$ then the overall rank of s must be higher than that of s'.

(2) If the average of the risk scores of student s is equal to that of student s', then the student with the higher standard deviation across their risk scores will have a higher overall score. This criterion would give a higher overall ranking score to students that have a high score on some risk factors over those that have a medium score across all risk factors.

In our algorithm we use the function, $overallRank_s : [0, 1] \rightarrow [0, 1]$, $overallRank_s = \sum_{i=1}^{L} o_i^{10}$ which satisfies both of the given criteria to aggregate outcomes.

3.2 Interface

The SIF's interface includes a two-level inspection panel that allows users to select a set of engagement and assessment metrics and visualises student data based on the resulting ranking. The first-level panel (see Figure 1) shows student information cards, ordered by the overall ranking scores according to three levels of priority (high, medium, and low; color-coded in red, yellow and gray). Each card provides a high-level overview of the student's learning behaviour, including (1) threshold flags for the risk factors with a high score (e.g., red indicating scores > 90), and (2) visualisation of the student's engagement and performance in the selected metrics with a globally and locally comparable visualisation method, namely a radar chart that includes the student's data versus the first, second and third quartiles of the class. Each card can be expanded to show additional data (see Figure 2) on the student, including (1) demographic data, and (2) box plots providing globally comparable visualisation of student's scores for each of the five risk factors. Furthermore, a help interface is designed which provides information on how students have been ranked and guidelines on how to interpret the visualisations and the values.

3.3 Application

This section demonstrates an application of SIF using real data captured from an undergraduate-level course with 407 students delivered at The University of Queensland.

Student dataset. We use student data from a course that was given in Semester 2 of 2021 to 407 undergraduate students. The course was taught during the COVID pandemic and most of the activities were performed online. As the learning behaviour metrics, we have selected a set of assessments and engagement that students performed until the mid-semester break (Week 8 of the semester), including the assignments {Assignment-1, Assignment-2}, and weekly engagement with Blackboard learning platform. The

Figure 1: SIF's student data inspection interface.

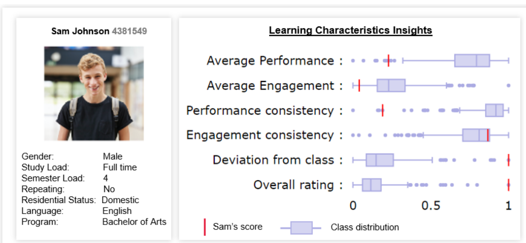

Figure 2: Detailed student information interface.

weekly engagement level is calculated by summing up the number of unique actions that each student takes in a day. Actions are deemed to be unique if they occur within a five-minute window.

Results. By applying SIF on the student data, students were ranked and categorised into high, medium, low, and no-priority for their likelihood of requiring the instructor's attention. We first explored the overall class behaviour in terms of the five characteristics and observed the following insights: as demonstrated by box plots in Figure 2 (in purple), (1) the average performance of students in the class is generally high, with outliers being at the bottom of the range, while their average weekly engagement is generally moderate, with outliers being at the top of the range. This difference is caused by assignment grades being capped while engagement numbers are uncapped; (2) most students are consistent in performance and engagement and have low class-deviation which work well for highlighting outliers with high deviation values, (3) overall ranking, also as per design, have a low value for the majority of the class, which leaves only a handful of cases for manual inspection.

Figure 1 demonstrates the overview cards for three students with high overall ranking scores, containing: 1) a radar chart presenting the input parameters (i.e., Assignment-1 and Assignment-2 scores, and weekly engagement level for the first eight weeks of the semester), 2) the flags raised for the student from the five risk factors. The main insights observed are: Student referred to as Sam Johnson, has low performance (at 23% of the class), and very low engagement (at 4% of the class), high inconsistency in performance (81%), and very high deviation from the class (at 100% of the class). Accordingly, in the radar chart, Sam's Assignment-2 is at the bottom of the class (affecting his average performance), while his Assignment-1 is at the second quartile of the class (making him inconsistent in performance). His engagement is below the first quartile of the class in all the weeks except in week-4, and the high deviation from the class might be explained by the low engagement in combination with a relatively high grade in assignment-2. Meng Fang, two flags have been raised indicating a high inconsistency in engagement and high deviation from the class. Accordingly, the radar chart for this student shows an extreme change of engagement level in every second week up to Week 7. Furthermore, outlining the class in most of the weekly engagements can explain her high class-deviation. Despite the changes in engagement, Meng has a high and consistent performance which can be discarded from a need for an intervention by the instructor. Amir Isa is reported within the medium priority category with a very low performance, low engagement, and medium deviation from the class. Accordingly, the radar chart shows that Amir's assignments are much lower than the class's first quartile, while he had a consistent low engagement in most of the weeks. Furthermore, outlining the class negatively in both of the assignments may explain his medium class-deviation.

4 Conclusion

This paper aims to facilitate instructors' proactive support by developing an explainable LA approach, namely Students Inspection Facilitator (SIF), for instructor-facing dashboards that utilises students' digital footprints to identify individual students that may need instructors' attention based on a set of risk factors grounded in literature. To illustrate the practical application of SIF, we used a data set obtained from an undergraduate course with 407 students delivered at The University of Queensland. The primary findings from inspecting the recommended students by SIF suggest that SIF has the potential to be integrated into LADs at the institution level to instantly and meaningfully interpret students' learning behaviour and rank them for interventions while still reserving the judgment for the instructor to decide whether to intervene or not. Our future work aims to conduct user studies to understand how instructors may use this approach.

Acknowledgement. This research was supported partially by the Australian Government through the Australian Research Council's Industrial Transformation Training Centre for Information Resilience (CIRES) project number IC200100022.

References

[1] Solmaz Abdi, Hassan Khosravi, Shazia Sadiq, and Dragan Gasevic. 2020. Complementing educational recommender systems with open learner models. In *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge*. 360–365.

[2] Nora'ayu Ahmad Uzir, Dragan Gašević, Wannisa Matcha, Jelena Jovanović, and Abelardo Pardo. 2020. Analytics of time management strategies in a flipped classroom. *Journal of Computer Assisted Learning* 36, 1 (2020), 70–88.

[3] Gökhan Akçapınar, Arif Altun, and Petek Aşkar. 2019. Using learning analytics to develop early-warning system for at-risk students. *IntInternational Journal Educ. Tech. Higher Educ.* 16, 1 (2019), 1–20.

[4] Mennatallah Amer and Markus Goldstein. 2012. Nearest-neighbor and clustering based anomaly detection algorithms for rapidminer. In *Proc. of the 3rd RapidMiner Community Meeting and Conference (RCOMM 2012)*. 1–12.

[5] Natasha Arthars, Mollie Dollinger, Lorenzo Vigentini, Danny Y-T Liu, Elsuida Kondo, and Deborah M King. 2019. Empowering teachers to personalize learning support. In *Utilizing learning analytics to support study success*. 223–248.

[6] et al Atherton. 2017. Using learning analytics to assess student engagement and academic outcomes in open access enabling programmes. *Open Learning: The Journal of Open, Distance and e-Learning* (2017), 119–136.

[7] David Baneres, M Elena Rodríguez-Gonzalez, and Montse Serra. 2019. An early feedback prediction system for learners at-risk within a first-year higher education course. *IEEE Transactions on Learning Technologies* 12, 2 (2019), 249–263.

[8] Hamsa Bastani, Osbert Bastani, and Carolyn Kim. 2018. Interpreting predictive models for human-in-the-loop analytics. *arXiv preprint arXiv:1705.08504* (2018), 1–45.

[9] Yizhou Fan, Wannisa Matcha, Qiong Wang, Dragan Gašević, et al. 2021. Learning Analytics to Reveal Links Between Learning Design and Self-Regulated Learning. *International Journal of Artificial Intelligence in Education* (2021), 1–42.

[10] Ed Foster and Rebecca Siddle. 2020. The effectiveness of learning analytics for identifying at-risk students in higher education. *Assessment & Evaluation in Higher Education* 45, 6 (2020), 842–854.

[11] D Gasevic, S Dawson, and A Pardo. 2016. How do we start? State and directions of learning analytics adoption. International Council for Open and Distance Education.

[12] Mary Sarah-Jane Gregory and Jason Michael Lodge. 2015. Academic workload: the silent barrier to the implementation of technology-enhanced learning strategies in higher education. *Distance education* 36, 2 (2015), 210–230.

[13] David Gunning and David Aha. 2019. DARPA's explainable artificial intelligence (XAI) program. *AI Magazine* 40, 2 (2019), 44–58.

[14] Lauren C Hensley, Christopher A Wolters, Sungjun Won, and Anna C Brady. 2018. Academic probation, time management, and time use in a college success course. *Journal of College Reading and Learning* 48, 2 (2018), 105–123.

[15] et al Holmes. 2021. Ethics of AI in education: towards a community-wide framework. *Int.International Journal of Artificial Intelligence in Education* (2021), 1–23.

[16] Jelena Jovanović, Dragan Gašević, Shane Dawson, Abelardo Pardo, Negin Mirriahi, et al. 2017. Learning analytics to unveil learning strategies in a flipped classroom. *The Internet and Higher Education* 33, 4 (2017), 74–85.

[17] Jelena Jovanović, Dragan Gašević, Abelardo Pardo, Shane Dawson, and Alexander Whitelock-Wainwright. 2019. Introducing Meaning to Clicks: Towards Traced-Measures of Self-Efficacy and Cognitive Load. In *Proceedings of the 9th International Conference on Learning Analytics & Knowledge* (Tempe, AZ, USA) *(LAK19)*. Association for Computing Machinery, New York, NY, USA, 511–520. https://doi.org/10.1145/3303772.3303782

[18] Hassan Khosravi, Gianluca Demartini, Shazia Sadiq, and Dragan Gasevic. 2021. Charting the design and analytics agenda of learnersourcing systems. In *LAK21: 11th International Learning Analytics and Knowledge Conference*. 32–42.

[19] Hassan Khosravi, Shiva Shabaninejad, Aneesha Bakharia, Shazia Sadiq, Marta Indulska, and Dragan Gašević. 2021. Intelligent Learning Analytics Dashboards: Automated Drill-Down Recommendations to Support Teacher Data Exploration. *Journal of Learning Analytics* 8, 3 (2021), 133–154.

[20] Haim Levy and Myles Robinson. 2006. *Stochastic dominance: Investment decision making under uncertainty*. Vol. 34.

[21] Wannisa Matcha, Dragan Gašević, Abelardo Pardo, et al. 2019. A systematic review of empirical studies on learning analytics dashboards: A self-regulated learning perspective. *IEEE Transactions on Learning Technologies* 13, 2 (2019), 226–245.

[22] Christoph Molnar. 2020. *Interpretable machine learning*. Lulu. com.

[23] Abelardo Pardo et al. 2018. OnTask: Delivering data-informed, personalized learning support actions. *Journal of Learning Analytics* 5, 3 (2018), 235–249.

[24] Dan Pelleg, Andrew W Moore, et al. 2000. X-means: Extending k-means with efficient estimation of the number of clusters.. In *Icml*. 727–734.

[25] Luis P Prieto, María Jesús Rodríguez-Triana, Roberto Martínez-Maldonado, Yannis Dimitriadis, and Dragan Gašević. 2019. Orchestrating learning analytics (OrLA): Supporting inter-stakeholder communication about adoption of learning analytics at the classroom level. *Australasian Journal of Educational Technology* 35, 4 (2019).

[26] Irina Rets, Christothea Herodotou, Vaclav Bayer, Martin Hlosta, and Bart Rienties. 2021. Exploring critical factors of the perceived usefulness of a learning analytics dashboard for distance university students. *Int J. Educ. Tech. Higher Educ.* (2021), In–Press.

[27] Mohammed Saqr, Uno Fors, and Matti Tedre. 2017. How learning analytics can early predict under-achieving students in a blended medical education course. *Medical teacher* (2017), 757–767.

[28] Beat A Schwendimann, Maria Jesus Rodriguez-Triana, Andrii Vozniuk, Luis P Prieto, Mina Shirvani Boroujeni, Adrian Holzer, Denis Gillet, and Pierre Dillenbourg. 2016. Perceiving learning at a glance: A systematic literature review of learning dashboard research. *IEEE Transactions on Learning Technologies* 10, 1 (2016), 30–41.

[29] Shiva Shabaninejad, Hassan Khosravi, Marta Indulska, Aneesha Bakharia, and Pedro Isaias. 2020. Automated insightful drill-down recommendations for learning analytics dashboards. In *Int.Proceedings of the Tenth International Conference on Learning Analytics & Knowledge*. 41–46.

[30] Shiva Shabaninejad, Hassan Khosravi, Sander JJ Leemans, Shazia Sadiq, and Marta Indulska. 2020. Recommending insightful drill-downs based on learning processes for learning analytics dashboards. In *Int.International Conference on Artificial Intelligence in Education*. 486–499.

[31] Varshita Sher, Marek Hatala, and Dragan Gašević. 2020. Analyzing the consistency in within-activity learning patterns in blended learning. In *Proceedings of the Tenth Int.International Conference on Learning Analytics & Knowledge*. 1–10.

[32] Billy Tak-ming Wong and Kam Cheong Li. 2020. A review of learning analytics intervention in higher education (2011–2018). *Journal of Computers in Education* 7, 1 (2020), 7–28.

[33] Jianing Zhou and Suma Bhat. 2021. Modeling Consistency Using Engagement Patterns in Online Courses. In *LAK21: 11th International Learning Analytics and Knowledge Conference*. 226–236.