

Automated Insightful Drill-Down Recommendations for Learning Analytics Dashboards

Shiva Shabaninejad

s.shabaninejad@uq.edu.au

The University of Queensland

Hassan Khosravi

h.khosravi@uq.edu.au

The University of Queensland

Marta Indulska

m.indulska@uq.edu.au

The University of Queensland

Aneesha Bakaria

a.bakaria1@uq.edu.au

The University of Queensland

Pedro Isaias

pedroisaias@gmail.com

The University of Queensland

ABSTRACT

The big data revolution is an exciting opportunity for universities, which typically have rich and complex digital data on their learners. It has motivated many universities around the world to invest in the development and implementation of learning analytics dashboards (LADs). These dashboards commonly make use of interactive visualisation widgets to assist educators in understanding and making informed decisions about the learning process. A common operation in analytical dashboards is a ‘drill-down’, which in an educational setting allows users to explore the behaviour of sub-populations of learners by progressively adding filters. Nevertheless, drill-down challenges exist, which hamper the most effective use of the data, especially by users without a formal background in data analysis. Accordingly, in this paper, we address this problem by proposing an approach that recommends insightful drill-downs to LAD users. We present results from an application of our proposed approach using an existing LAD. A set of insightful drill-down criteria from a course with 875 students are explored and discussed.

CCS CONCEPTS

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

1 INTRODUCTION

The educational data revolution, empowered by the increasing use of technology in education, has enabled universities to collect rich data on their learners. To help make sense of data on learners, a relatively new field commonly referred to as learning analytics (LA) has emerged. LA aims to analyse data on learners for the purposes of understanding and improving student learning, taking into consideration the diverse needs of the student population [20]. Within LA, development and evaluation of learning analytics dashboards (LADs) have received significant attention [9, 15, 19].

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. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LAK '20, March 23–27, 2020, Frankfurt, Germany

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7712-6/20/03...\$15.00

<https://doi.org/10.1145/3375462.3375539>

Earlier models of LADs relied significantly on predictive analytics. For example, Course Signals [1] used predictive models to determine students’ risk status. However, with the wide-adoption of predictive models to support automatic decision-making, the study of discrimination and bias in algorithms has attracted a lot of attention, questioning the ethics of the grounds (such as gender and ethnicity) on which decisions are being made [8]. As a consequence, more recent LADs have generally focused on directly reporting data by relying solely on stakeholder judgement for understanding and interpreting the results [9]. However, as the volume, velocity, and variety of learner data increases, making sense of multi-dimensional data becomes more challenging [13].

A potential alternative approach to using predictive models for analysing large data sets is to employ operations from OLAP (Online Analytics Processing) [7]. In particular, using the OLAP drill-down operation enables a user to meaningfully zoom in to explore data at a more granular level. In the educational setting, a drill-down would be performed by progressively adding filters on learner attributes (e.g., ‘program = CS’ and ‘videos viewed < 5’) to explore the behaviour of sub-populations of learners. The filtering feature can be beneficial in acquiring knowledge about learning behaviour of sub-populations as well as seeking possible correlations between students’ features and their learning state. While the drill-down operation is very powerful in terms of providing additional insights, drilling down into a complex multi-dimensional data set is a challenging task [10, 13]. The challenges, which are further explored in detail in Section 3, include not knowing what attributes to use, failing to find insightful results and encountering a drill-down fallacy “where incomplete insights result from potentially confounding factors not explored along a drill-down” [12].

In this paper, we explore developing LADs that recommend insightful drill-down criteria to instructors. Our proposed approach defines the notion of an insightful drill-down as finding a set of filtering rules that identify a sub-population of learners that deviate from class norms based on an instructor provided metric. Accordingly, we introduce an algorithm, referred to as Automated Insightful Drill-Down (AID), that employs a decision tree and a scoring function based on Kullback-Leibler (KL) divergence to recommend a set of up to k most insightful drill-down criteria. Unlike predictive models that automatically label students, AID uses machine learning algorithms to guide data exploration while still reserving judgement about interpreting student learning to instructors.

We present an application of our proposed approach using an existing LAD called Course Insights that is equipped with manual

Table 1: A sample of a students' dataset

#	Residential Status	Engagement	Assessment Score	#	Residential Status	Engagement	Assessment Score
S1	Domestic	High	High	S12	International	High	Low
S2	Domestic	High	High	S13	International	High	Mid
S3	Domestic	High	High	S14	International	High	Mid
S4	Domestic	High	High	S15	International	High	Mid
S5	Domestic	High	High	S16	International	High	High
S6	Domestic	High	Mid	S17	International	Low	Low
S7	Domestic	Low	Low	S18	International	Low	Low
S8	Domestic	Low	Low	S19	International	Low	Low
S9	Domestic	Low	Low	S20	International	Low	Low
S10	Domestic	Low	Mid	S21	International	Low	Low
S11	International	High	Low	S22	International	Low	Mid

drill-down functionality. We present information on how this drill-down functionality has been used in a manual unaided manner by 71 teaching staff members who have been involved in teaching courses that used Course Insights. We then take the case of one such course, with 875 students, and explore how the manual drill-downs compare to filters recommended by AID. Initial findings suggest that AID has the potential to be effectively integrated into LADs to provide insightful and instant drill-down recommendations.

2 RELATED WORK

Learning Analytics Dashboards (LADs). Grounded in data-driven analytics [2], LADs are defined as “a single display that aggregate different indicators about learner(s), learning process(es) and/or learning context(s) into one or multiple visualization” [19]. Within LA, development and evaluation of LADs have received significant attention. Early work from the LA community had a strong emphasis on the development of predictive models [14]. Despite the high success rate of predictive analytics, the use of predictive LA dashboards are associated with ethical concerns and dilemmas that limits their impact and adoption [3, 5]. An alternative approach to guide data exploration without ethical concerns is to provide recommendations via smart drill-downs.

Smart Drill-down Approaches The idea of data-driven exploration of OLAP operations is not new and has been used since the 90s [18]. However, the use of these techniques in the context of providing smart drill-down in visualising dashboards is quite recent. In particular, Parameswaran and colleagues have recently articulated the need for developing smart drill-downs for effective, interactive data exploration [11, 12]. Many of the proposed methods for generating smart drill-downs focus on detecting anomalies and finding small data portions with the highest difference (e.g. [12, 16, 17]) while some others focus on finding interesting differences in larger data subsets (e.g. [11]). Our proposed approach, AID, enables the user of LADs to request drill-downs at a level of granularity they are interested in (and takes a similar approach to (e.g. [18])).

3 CHALLENGES WITH MANUAL DRILL-DOWNS

Although the filtering system in an analytic dashboard allows finding interesting insights about the data, drilling down into a highly multi-dimensional data set can be time-consuming and error prone. In this section, we describe the three main challenges that have been identified in previous studies. To contextualise these challenges for the context of LADs, we use an illustrative example that uses a student-based dataset shown in Table 1.

In this example, four attributes of *Student number*, *Residential Status*, *Video Engagement* and *Assessment Score* are captured for 22 students. Considering *Assessment Score* as the target feature and that *Residential Status* and *Video Engagement* each have two possible values in our data set, 8 drill-down actions can be performed. Figure 1, which is inspired by the example used in [12], shows all the possible drill-downs (labeled as B to I).

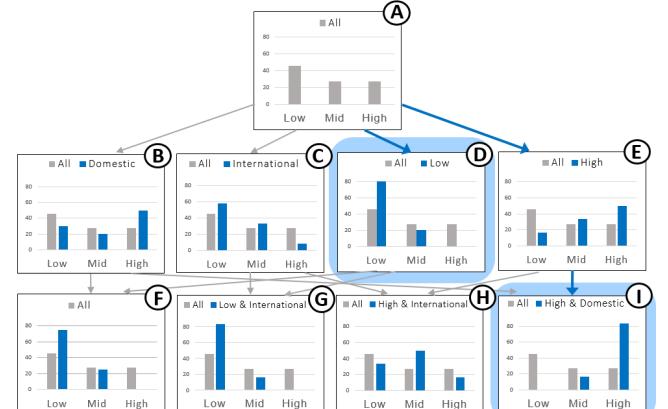


Figure 1: Drill-down illustration for the example given in Table 1. (D) and (I) represent insightful drill-downs.

Each chart represents assessment-scores’ distributions of the sub-population versus the overall population. Each bar in the charts represents the percentage of students with score value of High, Mid or Low. Two drill-down criteria marked by blue arrows identify two sub-populations D (students with Video Engagement=‘Low’) and I (students with Video Engagement=‘High’ and Residential Status=‘Domestic’) with a significant difference in their assessment result compared with the overall population.

Considering the presented drill-down criteria, three main challenges with a manual drill-down analysis may be encountered, which we discuss in the following.

Too many drill-down choices: To find interesting insights about sub-populations of students, the user has to determine what features to use for filtering. Trying all possible combinations of features may not be feasible. As the number of attributes and attribute values increases, the number of possible drill-down actions increases exponentially.

Lack of insightful results: A large number of possible drill-down criteria can lead to results which may not seem insightful. As a consequence, insightful drill-downs may be missed via a manual drill-down process. Referring to the example provided in Figure 1, many of the drill-downs, such as those presented by sub-populations B, C, H, may not seem insightful as the filtered population has similar assessment results compared to the entire population. In contrast, the two highlighted sub-populations D and I have significantly different assessment results compared with the overall population. In sub-population D 80% of students have a low score, which is almost 35% higher compared with the overall population. On the contrary, more than 83% of the sub-population I obtained a high score, which is 56% higher than the rate of high scores in the overall population.

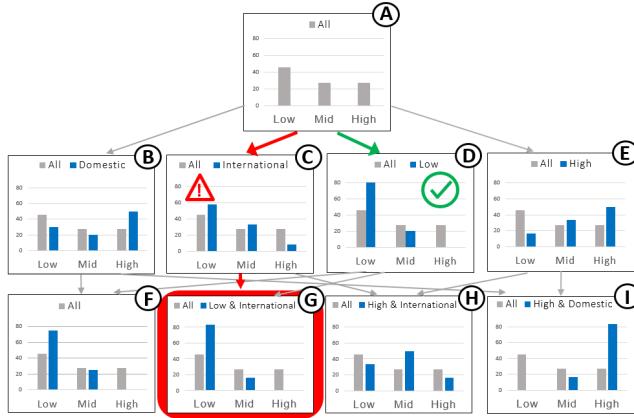


Figure 2: Drill-down fallacy illustration

Drill-down fallacies: An incomplete drill-down analysis may produce a class of errors known as *drill-down fallacy* [12]. Such a fallacy occurs when incorrect reasoning for a deviation found in the dataset is attributed to a smaller sub-population while “in fact it is a more general phenomenon” [12]. For instance, as shown in Figure 2, if the user takes the drill-down path that is marked by the red arrows, they could conclude that the high percentage of ‘Low’ scores and the absence of high scores (shown chart G) is associated to international students with ‘Low’ video engagement. However, if students are filtered first to ‘Low’ video engagement (the path represented in green), the user would understand that this distribution is associated to ‘Low’ video engagement in general and not international students with ‘Low’ video engagement specifically.

4 AUTOMATED INSIGHTFUL DRILL-DOWN (AID) RECOMMENDATIONS

In this section we introduce our proposed method for recommendation of insightful drill-down criteria in LADs. Section 4.1 introduces a formal notation and our problem statement. Section 4.2 presents our approach and section 4.3 presents an illustrative example based on the data set introduced in Section 3.

4.1 Notation and Problem Statement

Assume that D is a data set containing information about a set of students $S = \{s_1, \dots, s_N\}$ on a set of features $F = f_1, \dots, f_M$. Data pertaining to student s_n is represented in row n of D and data pertaining to feature m is represented in column m of D . Therefore, d_{nm} stores information on students s_n based on feature f_m .

A rule r is defined as assigning a value to a feature (e.g., ‘program’ = ‘ComputerScience’). For a feature with numerical values in D , the corresponding rule value can be a range instead of a single value (e.g., ‘age’ > 25). A drill-down criteria σ is defined as the conjunction of a set of rules (e.g., ‘program’ = ‘ComputerScience’ \wedge ‘age’ > 25). A drill-down criteria σ is said to cover a student s_n (i.e., include row n of D), if all of the rules in σ match the corresponding values for s_n . Consequently, applying σ to D leads to the selection of a set of students $S' \subseteq S$ such that σ covers each $s_n \in S'$. We define the coverage of a drill-down criteria C_σ as $\frac{|S'|}{|S|} \times 100$, which is the percentage of students S covered in the resulting sub-population S' . Based on this notation, our formal problem statement is as follows.

Formal problem statement: Given a data set D , a user provided target feature $f_t \in F$ and a constant $0 \leq \alpha \leq 1$, find a set of drill-down criteria $\Sigma = \{\sigma_1, \dots, \sigma_K\}$ that makes use of a set of features $F' \subseteq F - \{f_t\}$ such that each criteria σ_k : (1) selects a sub-population of learners S' that deviate significantly from the class norms on f_t and (2) the coverage of σ_k is bigger than α (i.e., $C_{\sigma_k} > \alpha$).

4.2 Proposed Approach

We present our approach, by first providing a high-level overview of the underlying algorithm, and then describing the automatic drill-down process using an example. We approach the problem of finding smart drill-downs that lead to the selection of sub-populations with significant deviation from the overall population as a classification algorithm to be solved using decision trees. We opt for this method because (1) decision trees are a well established method for classification with fast model construction and good model understandability; (2) a decision tree provides a tree-like model with each branch representing a set of feature-value pairs that can be mapped to a drill-down criteria; (3), the constant α can be given as a parameter to decision tree algorithms to limit the minimum size of classified sub-populations.

Algorithm 1 provides the high-level pseudo-code of our proposed approach. It takes five parameters as input: the dataset D , the features to be used in the drill-down criteria F' , the user provided target feature f_t , the minimum coverage by each recommended drill-down criteria α and the number of drill-down criteria to be recommended k . The output of the algorithm is a set of k drill-down criteria represented by Σ . The algorithm consists of three main blocks: create a decision tree; consider and score all tree paths; and find and return the top k drill-down criteria, as described in the remainder of this section.

Create a decision tree. The *LearnTree* function takes four parameters as input: the dataset D , the features to be used in the drill-down criteria F' , the target feature f_t and the minimum percentage of sub-population α . The output of this method is a decision tree T . *LearnTree* classifies the given dataset using a decision tree algorithm and generates a tree-like collection of nodes, where each node represents a splitting rule r for one feature, and each leaf (i.e. node with no children) provides the value distribution of the target feature for the sub-population S' . The tree is generated in such a way that the size of the sub-population in every leaf is at least equal to $\alpha\%$ of the overall population. *obtainTreePaths* takes the generated decision tree T as input and returns a set of drill-down paths P . Each path consists of a set of feature-value pairs as well as the distribution of the target feature values in the resulting sub-population.

Consider and score all tree paths. The decision tree embodies several possible drill-down paths, of which not all will necessarily result in a sub-population with a significantly different behaviour. *scorePath* takes each path p as input and scores the path by measuring the distance between S and S' using KL divergence defined as $\sum_{x \in X} S'(x) \log(\frac{S'(x)}{S(x)})$, where d is the distance between the distribution of the values in S and S' . And, X is a set that consists of the distribution of each target value in S' .

Find and return the top k drill-down criteria. Finally, our algorithm picks the k highest scored paths, converts them to a set of drill-down criteria σ and returns a recommendation to the user.

Algorithm 1: Finding a set of k smart drill-down criteria

```

Input :  $D, F', f_t, \alpha, k$ 
Output: $\Sigma$ 
1 create a decision tree;
2  $T \leftarrow LearnTree(D, F', f_t, \alpha);$ 
3  $P \leftarrow obtainTreePaths(T);$ 
4 Consider and score all tree paths;
5  $score[P] \leftarrow \{0\};$ 
6 foreach  $p \in P$  do
7   |  $score[p] \leftarrow scorePath(p);$ 
8 end
9 Find and return the top K drill-down criteria;
10  $\Sigma \leftarrow [];$ 
11  $topK \leftarrow topScores(score, k);$ 
12 foreach  $p \in topK$  do
13   | if  $score[p] > 0$  then
14     |   |  $\sigma \leftarrow turnPathToDrill-downCriteria(p);$ 
15     |   |  $\Sigma \leftarrow appendDrill-downCriteria(\sigma);$ 
16   | end
17 end
18 return  $\Sigma;$ 

```

4.3 Illustrative Example

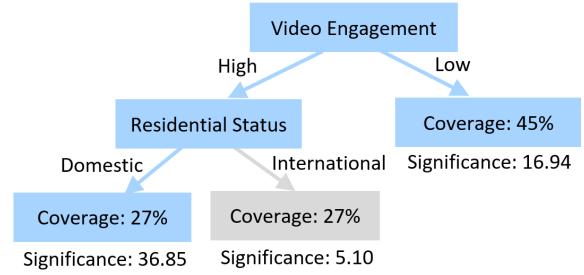
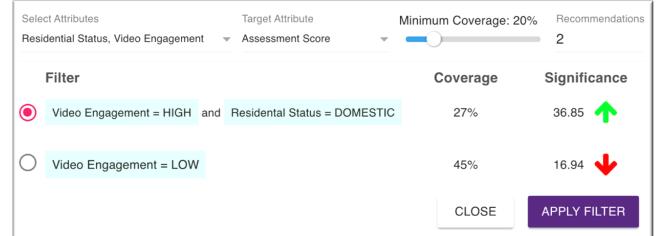
In this section, we explain how our algorithm is used to find the top 2 insightful drill-down criteria for the dataset D given in Table 1, {Residential Status, Engagement} as F' and Assessment-score as the target feature f_t . $k = 2$ and $\alpha=20\%$ are selected for this example to better demonstrate the approach. Initially, the dataset is classified using a decision tree algorithm. As shown in Figure 3, the resulting decision tree includes three leaves that can be mapped to drill-down paths of P_1 : [Video Engagement = ‘High’ and Residential Status=‘Domestic’], P_2 : [Video Engagement = ‘High’ and Residential Status=‘International’] and P_3 : [Video Engagement = ‘Low’].

Our proposed algorithm uses a scoring function based on KL divergence to evaluate the ‘insightfullness’ of each of these paths. For instance, path P_1 with the distributions: Low: 0%, Mid: 16.66%, High: 83.33%, and the distributions in the overall population: Low: 45.45%, Mid: 27.27%, High: 27.27% will be scored as $0 \log(\frac{0}{45.45}) + 16.66 \log(\frac{16.66}{27.27}) + 83.33 \log(\frac{83.33}{27.27}) = 36.85$. The score of P_2 (5.10) and P_3 (16.94) are computed in a similar way. These drill-down paths and their resulting coverage and significance scores are shown in Figure 1 (P_1 shown in Chart I, P_2 shown in Chart H, and P_3 shown in Chart D).

The two paths with the highest scores (P_1 and P_3) are converted to a set of drill-down criteria and are recommended to the user. Figure 4 shows the prototype of AID interface in Course Insights, representing the input and the resulting recommendations including the drill-down criteria, coverage and significance. Furthermore, an arrow indicating the average result of Assessment Score in the sub-population in comparison with the average result in the overall population. Note that the returned result satisfied the second criteria of the problem statement; the coverage of both drill-down criteria exceeds α .

5 APPLICATION

This section presents an application of AID using an existing LAD called Course Insights, which is equipped with manual drill-down functionality. Section 5.1 introduces Course Insights and its main segments. Section 5.2 provides insights on how manual drill-down functionality was employed by 71 Course Insights users. Finally, Section 5.3 presents results from applying our approach on student

**Figure 3: The resulting decision tree.****Figure 4: The drill-down criteria recommendation interface.**

data from one of the courses that have previously employed Course Insights. The recommended drill-down criteria and how they are presented to users are explored and discussed in this section.

5.1 Course Insights

Course Insights is a LAD that provides filterable and comparative visualisations of aggregated students’ daily activity and learning events. Course Insights aims to provide actionable insights for teachers by linking data from several sources including the Students Information System, Blackboard, edX Edge, and a number of embedded learning tools (such as Echo360, Kaltura) to create a multi-dimensional educational data set. Course Insights is equipped with a filtering system to enable instructors to drill-down into the data and explore the behaviour of sub-populations of learners.

Course Insights consists of seven analytical segments (i.e dashboard screens). Only the segments with filter related functionality is discussed here due to lack of space. (1) *Enrolment* providing statistical information about students’ demographic and educational features; (2) *Engagement* providing time-based statistical data about online learning activities; (3) *Assessment* providing statistical data about students’ assessment results and (4) *Students* providing intervention facility enabling educators to use filters to identify and contact groups of students. The Students segment, displays a tabular view with a row for each student and columns for all features collated for a student. A Prototype of the platform demonstrating the instructor view is available online at <https://analytics.itali.uq.edu.au/dev/insights/index.html>.

Figure 5 illustrates the two filter interfaces that are supported by Course Insights. The top part of the figure shows the Basic Filtering interface, which only allows drill-downs on students’ age, program, gender, study-load, and residential-status attributes. The bottom part of the figure shows the Advanced Filtering interface, which allows users to build their filters more freely. Advanced filtering allows users to select attributes from demographic, assessment, engagement and enrolment features. The selected set of attributes

Figure 5: Course Insights Filtering Interface



Figure 6: Statistical and graphical demonstration of filtered versus unfiltered students' results for a target feature.

can be combined using AND (presented by ‘Match All’) or OR (presented by ‘Match Any’). When a filter is applied in Course Insights, statistical data as well as a graph representing the filtered versus unfiltered distribution of the target feature (as demonstrated in Figure 6) are presented to the user.

5.2 Insights: Manual Drill-downs

To understand the common challenges with manual drill-downs in LADs, and to study how they are being used, we have analysed logs from Course Insights. Our investigation was guided by the following questions: (1) how complex are the manual drill-downs applied by users? and (2) which manual drill-downs are commonly applied? To answer these questions, we processed 356 manual drill-down actions performed by 71 teaching staff members who were involved in teaching courses during 10 weeks of a semester at The University of Queensland in 2019. To answer the first question, the average length of each drill-down action was examined. To answer the second question, the frequency of the use of performed manual drill-down actions was examined. At the time the data was collected, the Students segment (i.e. dashboard screen) was the only segment that provided advanced filtering.

Analysis of data related to the first question shows that using a combination of attributes was less common. Overall, 84% of drill-downs were performed on a single attribute and the average number of attributes for each drill-down was <2 (Enrolment = 1.35 ± 0.55 , Engagement = 1.34 ± 0.68 , Assessment = 1.29 ± 0.59

Table 2: Top 10 common manual drill-downs

Single Criteria	total	Multi Criteria	total
Age	98	Age & Study Load	22
Program	71	Age & Gender	6
Residential Status	66	Program & Gender	6
Study Load	35	Program & Age	6
Gender	32	Program & Residential Status & Study Load	5

and Students = 1.15 ± 0.36). Interestingly, despite having access to the Advanced Filter, the Student segment had the lowest average length of drill-down actions. This result might be explained by the well-established phenomena called *choice overload* [6] from Choice Theory that suggests increasing the number of choices given to a user may increase their effort in decision making. The 10 most common drill-down actions are reported in Table 2. The results of this table confirm the findings from the previous study, indicating that most of the drill-downs were consist of a single-attribute. Only one of the 10 most common drill-downs used a combination of more than two attributes.

As evident, teaching staff have commonly used a small subset of the features and rarely drilled-down more than one level into the features. This may suggest that the drill-down functionality has been under-utilised and that many users may require assistance to fully benefit from “drill-down” operation. An alternative explanation supported by findings of Wise and Yeonji [21] is that “instructors did not always come to analytics use with specific questions, but rather with general areas of curiosity”. Examples of this can include questions such as ‘how did international students or female students perform compared to other students, which were commonly asked, as evident by the results of Table 2, by Course Insights’s users. Our future work aims to partner with academics through a co-creation lens [4] to investigate (1) how manual drill-down criteria are formed by academics, (2) how academics judge the insightfulness of a given drill-down operation, and (3) how academics compare and contrast data-driven exploration compared to curiosity-driven exploration in LADs.

5.3 Insightful Drill-Downs in Action

This section takes the case of one of the courses that have used Course Insights to explore how the manual drill-downs compare to recommendations provided by AID. For this comparison, we use an introductory course on ‘Calculus and linear algebra’ offered in 2019 to 875 undergraduate students. We have selected a set of attributes: {Brand New (YES, NO), edX Edge click-stream (High, Low, Med), Gender (Male, Female, Other), Program (48 different programs), Residential Status (Domestic, International)} as the selected attributes F' and the assessment called ‘final exam-Section A’ as the target feature f_t . A total of 1728 drill-down actions were possible for this data set. The reported results are based on running AID three times for three different α : 2%, 10% and 30%, and the value $k = 3$ as the number of most significant drill-down criteria. The three variations of α are used to demonstrate its impact on the recommendations. In all reported cases, the recommendations were generated instantly.

Figure 3 presents the recommendations generated for this course. The arrows indicate the average result of the target attribute in the sub-population compared with the average result in the overall population as in Figure 4. As demonstrated, the use of $\alpha = 30\%$ leads to

Table 3: Resulting recommendations generated by AID.

α	Recommended Drill-Down Criteria	Coverage	Significance
2%	(1) [Program = Bachelor of Computer Science]	6%	19.15 ↑
	(2) [Program = Bachelor of Engineering (Honours) and Master of Engineering]	8%	8.90 ↑
	(3) [Program = Bachelor of Engineering (Honours)] and [Residential status = International] and [edX Edge engagement = Mid]	2%	8.44 ↑
10%	(4) [Gender = MALE] and [Residential status = Domestic] and [edX Edge engagement = Low]	21%	1.70 ↓
	(5) [Gender = Female]	15%	1.24 ↑
	(6) [Gender = Male] and [Residential status = Domestic] and [EdX Edge engagement = High]	15%	1.19 ↓
30%	(7) [edX Edge engagement = Low]	33%	0.70 ↓
	(8) [edX Edge engagement = High]	33%	0.46 ↑
	(9) [edX Edge engagement = Mid]	33%	0.04 ↑

the generation of simple drill-down actions (average length = 1) with a low significance score (average significance = 0.40) that find relatively large sub-populations (avg coverage = 33%). In comparison, the use of $\alpha = 10\%$ leads to the generation of more complex drill-down actions (average length = 2.33) with a higher significance score (average significance = 1.38) that find smaller sub-populations (average coverage = 17%). The generated recommendations for $\alpha = 2\%$ provide some interesting insights: compared to the previous two cases, as expected the significance score has increased (average significance = 12.28), and the sub-populations size has decreased (average coverage = 5.3); however compared to $\alpha = 10\%$, the average length of the drill-down criteria has decreased (1.6).

The main insights gained from our application of AID are (1) AID can generate insightful (high significance score) drill-down recommendations instantly, making it suitable for integration into LADs (2) Contrary to our initial assumptions, the average length of a drill-down action is not necessarily a good indicator of its insightfulness. However, it is worth noting that finding insightful drill-down criteria of even size 1 in the presence of attributes with a large number of potential values (e.g. program with 48 different values) is still a challenging task; (3) α plays a very important role in the quality of the generated recommendations. A poor selection in α may lead to disappointing results for the user. Our future work aims to (1) extend AID such that the algorithm itself suggests promising values of α as well as to partner with academics through a co-creation lens to investigate (2) what criteria (perhaps instead of maximising deviation) can be used for determining insightfulness and (3) what would be the best way, if at all, to present recommendation drill-down criteria to academics.

6 CONCLUSION

The OLAP drill-down operation is commonly used in data-driven dashboards to enable users to meaningfully zoom in to explore data at a more granular level. In the context of LADs, this operation can be used to enable users to identify a sub-population of learners who deviate from class norms and who may require special attention. In this paper, we outlined common challenges with manual drill-downs in LADs and address these challenges by providing a method, referred to as AID for finding and recommending a set of insightful drill-down actions. AID uses a decision tree classification method

to examine the insightfulness of all drill-down paths and uses KL-divergence to rank and select the highest ranked paths for the user to consider.

Initial findings, based on analysis of user logs from a LAD equipped with manual drill-down functionality, suggest that manual drill-downs may be under-utilised by LAD users. Using AID to generate recommendations for a course that has used this LAD demonstrates that AID has the potential to be effectively integrated into LADs to provide insightful (with high significance score) and instant drill-down recommendations to guide data exploration while still reserving judgement about interpreting student learning to instructors. In future research, we aim to partner with LAD users to explore the practical implications of AID in LADs and refine it based on the derived insights.

REFERENCES

- [1] Kimberly E. Arnold and Matthew D. Pistilli. 2012. Course Signals at Purdue: Using Learning Analytics to Increase Student Success. In *ICLAK*. 267–270.
- [2] Robert Bodily, Judy Kay, Vincent Aleven, Ioana Jivet, Dan Davis, Franceska Xhakaj, and Katrien Verbert. 2018. Open Learner Models and Learning Analytics Dashboards: A Systematic Review. In *ICLAK*. 41–50.
- [3] Linda Corrin, Gregor Kennedy, et al. 2019. The ethics of learning analytics in Australian higher education.
- [4] Mollie Dollinger, Danny Liu, Natasha Arthars, and Jason M Lodge. 2019. Working Together in Learning Analytics towards the Co-Creation of Value. *Journal of Learning Analytics* 6, 2 (2019), 10–26.
- [5] Hendrik Drachsler, Tore Hoel, Maren Scheffel, et al. 2015. Ethical and Privacy Issues in the Application of Learning Analytics. In *ICLAK*. 390–391.
- [6] John T Gourville and Dilip Soman. 2005. Overchoice and assortment type: When and why variety backfires. *Marketing science* (2005), 382–395.
- [7] Jim Gray, Surajit Chaudhuri, et al. 1997. Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals. *DMKD* (1997), 29–53.
- [8] Sara Hajian, Francesco Bonchi, and Carlos Castillo. 2016. Algorithmic Bias: From Discrimination Discovery to Fairness-aware Data Mining. In *ACM SIGKDD ICKDDM*. 2125–2126.
- [9] Ioana Jivet, Maren Scheffel, Marcus Specht, and Hendrik Drachsler. 2018. License to Evaluate: Preparing Learning Analytics Dashboards for Educational Practice. In *ICLAK*. 31–40.
- [10] Manas Joglekar, Hector Garcia-Molina, and Aditya Parameswaran. 2015. Smart Drill-down: A New Data Exploration Operator. *PVLDB* (2015), 1928–1931.
- [11] Manas Joglekar, Hector Garcia-Molina, and Aditya Parameswaran. 2017. Interactive data exploration with smart drill-down. *IEEE Transactions on Knowledge and Data Engineering* 31, 1 (2017), 46–60.
- [12] Doris Jung-Lin Lee, Himel Dev, Huizi Hu, Hazem Elmeleegy, and Aditya Parameswaran. 2019. Avoiding Drill-down Fallacies with VisPilot: Assisted Exploration of Data Subsets. In *ICII*. 186–196.
- [13] Doris Jung-Lin Lee and Aditya Parameswaran. 2018. The Case for a Visual Discovery Assistant: A Holistic Solution for Accelerating Visual Data Exploration. *Data Engineering* (2018), 3–14.
- [14] Farshid Marbouti, Heidi A Diefes-Dux, and Krishna Madhavan. 2016. Models for early prediction of at-risk students in a course using standards-based grading. *Computers & Education* (2016), 1–15.
- [15] Wannisa Matcha, Nora'Ayu Ahmad Uzir, Dragan Gasevic, and Abelardo Pardo. 2019. A Systematic Review of Empirical Studies on Learning Analytics Dashboards: A Self-Regulated Learning Perspective. *IEEE TLT* (2019), 1–1.
- [16] Sunita Sarawagi. 2000. User-adaptive exploration of multidimensional data. In *VLDB*. 307–316.
- [17] Sunita Sarawagi. 2001. User-cognizant multidimensional analysis. *VLDB J.* (2001), 224–239.
- [18] Sunita Sarawagi, Rakesh Agrawal, and Nimrod Megiddo. 1998. Discovery-Driven Exploration of OLAP Data Cubes. In *ADT EDBT'98*. Springer, 168–182.
- [19] Beat A. Schwendimann, Maria Jesus Rodriguez-Triana, et al. 2017. Perceiving Learning at a Glance: A Systematic Literature Review of Learning Dashboard Research. *IEEE TLT* (2017), 30–41.
- [20] George Siemens and Dragan Gasevic. 2012. Guest editorial-learning and knowledge analytics. *Journal of Educational Technology & Society* 15, 3 (2012), 1–2.
- [21] Alyssa Friend Wise and Yeonji Jung. 2019. Teaching with Analytics: Towards a Situated Model of Instructional Decision-Making. *Journal of Learning Analytics* 6, 2 (2019), 53–69.