

Design of a Course Recommender System as an Application of Collecting Graduating Attributes

by

Behdad Bakhshinategh

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

Department of Computing Science

University of Alberta

© Behdad Bakhshinategh, 2016

Abstract

In educational research, the term of Graduating Attributes has been used for the qualities, skills and understandings a university community agrees its students would develop. Having a description of Graduating Attributes is one of the ways through which universities can display the outcomes of higher education. But can Graduating Attributes be used also to enhance the process of learning? In this thesis, we discuss how graduating attributes can be used in data mining applications to improve the learning process. An example of a data mining application can be a course recommender system which helps students to choose the courses they would participate in. In our work we have implemented this recommender system as an example of possible applications which Graduating Attributes can provide. In order to achieve such a goal we first needed to implement a tool for assessing Graduating Attributes and gather data. In spite of implementing this tool, we were not able to gather sufficient amount of data. As a result, based on the structure of data in our assessment tool, we have generated synthetic data which we have used for the evaluation of the course recommender system. The results of the recommendation improve over time as a result of having more data. The mean squared error decreases from 0.32 in second semester to 0.08 in the tenth semester.

To the Count

For teaching me everything I need to know about math.

The good news about computers is that they do what you tell them to do.

The bad news is that they do what you tell them to do.

– Ted Nelson

Table of Contents

1	Introduction	1
1.1	Thesis Statement	2
1.2	Thesis Contributions	2
2	Educational Data Mining	4
2.1	End users in EDM tasks	5
2.2	Methods and Applications	6
2.3	Student modelling	8
2.3.1	Predicting Student Performance, Achievement of Learning Outcomes or Characteristics	9
2.3.2	Detecting Undesirable Student behaviours	10
2.3.3	Profiling and Grouping Students	10
2.3.4	Social Network Analysis	11
2.4	Decision Support Systems	11
2.4.1	Providing Reports	12
2.4.2	Creating Alerts for Stakeholders	12
2.4.3	Planning and Scheduling	13
2.4.4	Creating Courseware	13
2.4.5	Developing Concept Maps	13
2.4.6	Generating Recommendation	14
2.5	Other Applications	14
2.5.1	Adaptive Systems	14
2.5.2	Evaluation	15
2.5.3	Scientific Inquiry	15
2.6	Literature Review	16
2.7	Brief Summary	20
3	Graduating Attributes: Definition and Implementation	22
3.1	Competency Based Education	22
3.1.1	Design Criteria of Competencies	23
3.1.2	Assessment	25
3.1.3	Self-assessment	26
3.2	Graduating Attributes	27
3.2.1	Scales with intuitive labels (GA Levels of Acquisition)	28
3.2.2	Generic attributes Interpreted by a series of sub-attributes:	28
3.2.3	Appropriate description of acquisition levels with regards to KSA GAs:	29
3.3	Online Assessment Tool	29
3.3.1	Online Questionnaire	31
3.3.2	Student's Graduate Attribute Assessment Tool (GAAT)	31
3.3.3	Instructors Graduate Attribute Assessment Tool	35
3.3.4	Course Report for Instructors	36

3.4	Brief Summary	37
4	EDM Applications of Graduate Attribute Assessment Tool	38
4.1	Introduction	38
4.2	Matching Employees and Employers	40
4.2.1	Summary of Design	40
4.3	Course Recommender System	41
5	Course Recommender System	43
5.1	Introduction	43
5.2	Data	44
5.2.1	Structure of Data	44
5.2.2	Data Generation	45
5.3	Algorithm Design	49
5.4	Experiments and Evaluation	52
5.5	Future Work	55
6	Summary and Conclusions	57
	Bibliography	59
	Appendix A List of information targeted by the Online Questionnaire	64

List of Tables

2.1	Targeted Users of EDM Applications	7
3.1	Graduating Attributes and sub-attributes	29
5.1	Example of data representation: courses x competences x students	50

List of Figures

2.1	Taxonomy of Applications in EDM	8
2.2	Comparison of Reported applications in EDM	18
3.1	Student's Graduate Attribute Assessment Tool Snapshot . . .	32
3.2	Student's GAAT Snapshot showing attributes, sub-attributes and attribute descriptions	33
3.3	Student's GAAT Snapshot showing the contributed courses . .	34
3.4	Instructors Graduate Attribute Assessment Tool Snapshot . .	35
3.5	Course Report for Instructors	36
4.1	Example of matching students to job positions	41
5.1	Relations of data entities	45
5.2	Values assigned to some example courses in the generated data	47
5.3	Values assigned to an example sub-attribute over time for some example students in the generated data	48
5.4	Average ratio of ratings for all student over time	49
5.5	Mean squared error of rating predictor algorithms in each semester	53
5.6	Results of correlation thresholding for prediction at semester 7	54
5.7	Results of correlation thresholding for prediction at semester 9	54
5.8	Results of K best neighbours method on prediction	56

Acronyms

CBE Comptency Based Education.

CF Collaborative Filtering.

EDM Educational Data Mining.

GA Graduate Attribute.

GAAT Graduate Attribute Assessment Tool.

ITS Intelligent Tutoring System.

JEDM Journal of Educational Data Mining.

KSA Knowledge, Skills and Abilitie.

MOOC Massive Open Online Courses.

MSE Mean Squared Error.

Chapter 1

Introduction

In today's institutions of higher education, one of the topics of research is about assessing the achievement of intended learning outcomes of the students. An example of attention to this topic was the report published by University of Alberta's provost office in 2011, in which Graduating Attributes are proposed to describe the qualities, values and dispositions that students are developing [1]. One of the most commonly used definitions of Graduating attributes (GAs) suggested by Bowden et al. is:

“Graduating attributes are the qualities, skills and understandings a university community agrees its students should develop during their time with the institution. These attributes include but go beyond the disciplinary expertise or technical knowledge that has traditionally formed the core of most university courses. They are qualities that also prepare graduates as agents of social good in an unknown future.” [8]

Later on, in research done by Ipperciel and ElAtia, an assessment model was suggested for assessing the GAs [25]. This model introduces various values for assessing each of the Graduating Attributes as well as descriptions which can be used in all fields of higher education. This helps with providing a general understanding of the GAs for different departments.

In our work we have adopted the assessment model of GAs and developed an online assessment tool for Graduating Attributes, in order to create a platform for enhancing the usage of Graduating Attributes. This tool will help us with gathering data about about Graduating Attributes, which would enable

us with using them in data mining applications. This tool can also show us how different stakeholders would interact with a system targeting Graduating Attributes. In the rest of this thesis, we will mention this tool as Graduate Attribute Assessment Tool (GAAT).

1.1 Thesis Statement

The statement we tried to prove in this thesis is that Graduating Attributes can be used in developing new applications which can enhance the process of learning. Various applications can target various stakeholders involved in our system. Useful applications for each of the stakeholders can also help with the process of engaging them with the GAAT, by encouraging them to use the system. The different stakeholders we will mainly discuss in this thesis are learners, instructors, administrators and researchers. However we will mention applications which can be useful to stakeholders beyond the educational environment such as industries which are willing to employ graduated students.

The application we will implement as an example of all possible applications of Graduating Attributes is a course recommender system for learners.

1.2 Thesis Contributions

To reach our purpose, the first question we need to ask is how Graduating Attributes Assessment Tool can be implemented. The design should consider the different aspects of the data we intend to gather and how users would engage with the system. Users expect to receive a service in return for the time they may spend with the system and we need to recognise the possible services this tool can provide to each of the stakeholders. As a result, the first contribution of this thesis is design and implementation of an online assessment tool for assessing Graduating attributes and providing reports to stakeholders.

Then we need to know about the possible data mining applications of collecting data related to Graduating Attributes. To recognise these possible applications we will look into the different applications of educational data mining. Our contribution in this topic is a taxonomy of applications in educa-

tional data mining which we would also use for recognizing various application of Graduate Attribute Assessment Tool (GAAT).

Finally we need to look into one of the possible data mining applications of GAAT, namely generating course recommendation. Although we implemented the GAAT for gathering the data, we were not able to gather sufficient amount of data for the evaluation of a course recommender system. As a result, for training and testing our model we needed to generate synthetic data based on the structure of the data GAAT would gather. We will use this synthetic data for evaluation of our course recommender system.

Hence, in this thesis, we will first illustrate the applications of Educational Data Mining (EDM) in Chapter 2, in order to propose a taxonomy of tasks and applications in EDM. Based on the taxonomy, the possible applications of GAAT can be proposed. Then we will introduce the Graduating Attributes and how we have designed the online assessment tool in Chapter 3. This will help us with understanding the type of data which can be gathered by this system. In Chapter 4, we will introduce some of the possible applications of Graduate Attribute Assessment Tool, and then in Chapter 5, we will discuss generating synthetic data, our recommender system and the results of our experiments.

Chapter 2

Educational Data Mining

Educational Data Mining (EDM) is the field of using data mining techniques in educational environments. Currently there are many computer-based learning systems which gather large amounts of usage data such as Learning Management Systems (LMS), Massive Open Online Courses(MOOCs) and Intelligent Tutoring Systems (ITS). Stone et al. define an LMS as “a centralized web based information systems where the learning content is managed and learning activities are organized. LMS represents a more general term for a technology framework that supports all aspects of formal and informal learning processes, including learning management, content management, course management, etc.”[53]. It consolidates in one platform a number of functionalities, including “personal communication (email and instant messages), group communication (chatting and forums), content posting (syllabus, papers, presentations, lesson summaries), performance evaluation (question and answer repositories, self-assessment tests, assignments, quizzes and exams), and instruction management (message and grade posting, surveys, and online office hours)”[34] while serving as point of departure to the entire web. The data collected by LMSs provide us with the opportunity of using data mining.

There exist various methods and applications in EDM. These applications can follow both applied research objectives such as improving and enhancing learning quality, as well as pure research objectives, which tend to improve our understanding of the learning process. Aside from the classification of applications based on their objectives, which is the focus of this chapter, EDM

applications can also be categorized based on the targeted end user. The applications of EDM can target any of the stakeholders involved in educational systems, such as learners, educators, administrators and researchers themselves. Providing feedback, personalization and recommendations can improve the learning process of students. Making discoveries and providing decision support systems can help the educators by improving teaching performance and making decisions. Administrators are provided resources and tools for making decisions and organizing the institutions.

In this chapter, we will look into various possible applications of EDM and cognate methods that can be used to meet the needs of these applications. We focus on the end objectives of these applications and seek to group the applications in categories with similar purposes. Although we try to draw a line between different categories of applications, it should be noted that in some cases, there is no clear boundary between the applications. Some research may belong to more than just one of these categories. In some other cases, an application can be used as a tool for reaching another application. There are many such examples; for instance, creating reports of students' expected performance for educators. In this case, the end objective is providing reports which need some visualization techniques. However, prediction of student performance can be described as another application in EDM, which is needed prior to providing a report.

In the next section we will briefly discuss classification of EDM tasks based on the end user. Then we will introduce taxonomy of applications and tasks in EDM based on their end objectives. And finally we will look into the related studies and will compare them with our study.

2.1 End users in EDM tasks

For better clarification of the identified applications, we can look into the target users of each application. This has the added value of also showing the possible applications for the end users which have not been targeted yet. The end users in educational environments are learners (students), educators

(instructors), administrators and researchers.

Learners have been the target of EDM in various applications such as grouping students, generating recommendations and adaptive systems. One important goal of EDM as a whole is improving the quality of learning; and in the process of learning, two groups of users come to mind first, i.e. learners and educators. Most of the applications in categories of student modelling and decision support systems target educators as their end users. Student modelling provides a better understanding of students' state of learning and decision support systems can directly help educators make better decisions for improving the learning process. This also applies to the administrator of educational institutions making higher level decisions. Researchers also represent a group of end users, as the objective of the research is to understand the learning process, develop theories and test them. As an example, researchers can use social network analysis (SNA) to pinpoint the properties that are valuable in prediction of student performance. Table 2.1 presents possible targeted users of each application. It is important to mention that any research in EDM may address one or more than one of these classes at once.

2.2 Methods and Applications

The methods of EDM are the same as those in the data mining field in general. There are multiple methods to use in EDM for each of the various applications. Among these methods, the most used are classification and regression (1), clustering (2), association rule mining (3), discovery with models (4), outlier detection (5), social network analysis (6), text mining (7), sequential pattern mining (8) and visualization techniques (9) also referred to as distillation of data for human judgment. Based on research [46] by Romero and Ventura in 2009, the most commonly applied data mining tasks are regression, clustering, classification and association rule mining.

Applications and tasks in EDM can be categorized based on different properties. Multiple surveys of EDM exist, which have listed possible applications of EDM. We will look into these surveys in more detail in the literature review.

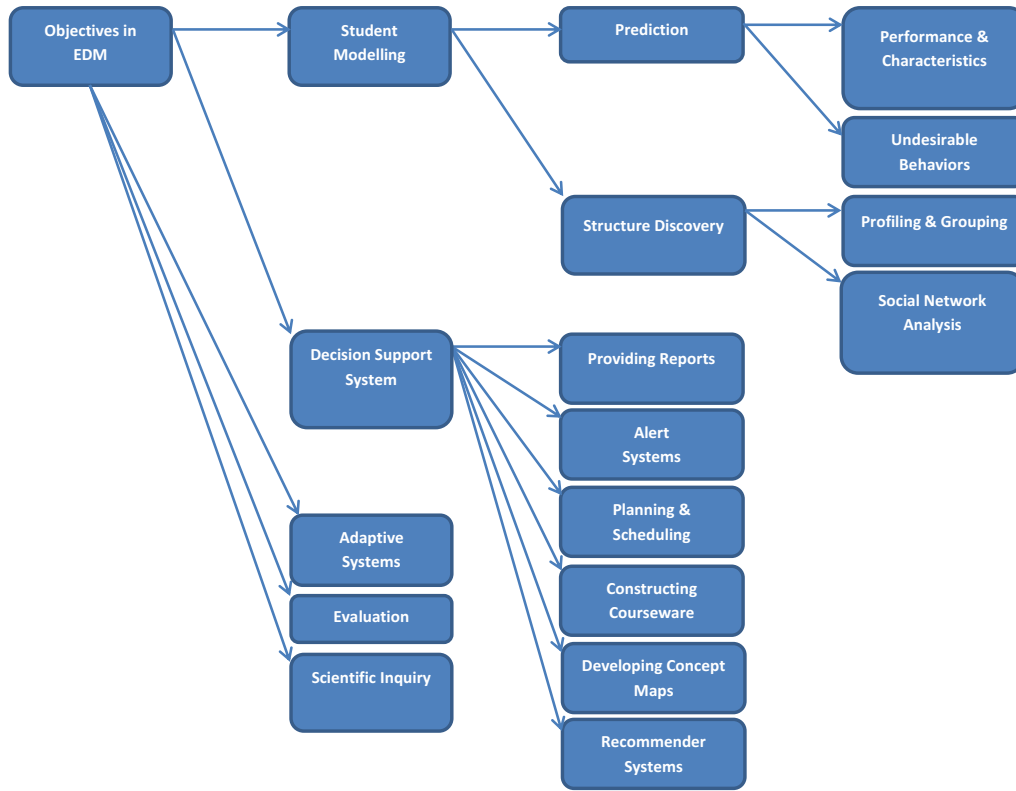
	Students	Educators	Administrators	Researchers
Predicting performance and characteristics		X	X	
Detecting undesirable student behaviour		X	X	
Profiling and Grouping students	X	X		
Social Network Analysis		X	X	X
Providing reports	X	X	X	
Creating alerts for stakeholders		X	X	
Planning and scheduling	X	X	X	
Constructing courseware		X		
Developing concept maps		X		X
Generating recommendation	X	X		
Evaluation		X		
Adaptive systems	X			
Scientific inquiry				X

Table 2.1: Targeted Users of EDM Applications

Considering these surveys and reading into the research examples they have provided, as well as the recent studies published in journals of educational data mining, we propose a new list of EDM categories. In this list, we have tried to consider all the categories mentioned in previous surveys and in the literature, as well as new categories which we think need to be added. These new categories of applications can be explained by the growth of interest in EDM. In previous studies, possible applications of EDM have been introduced sometimes in no specific order, sometimes based on the number of research papers done in each of the categories. We try to group possible applications of EDM into categories based on their end objective. We have tried to group different applications together as much as possible to better highlight the similarities and differences.

We have identified 13 categories of applications, as shown in Figure 2.1, forming a new taxonomy tailored specifically to EDM, thus setting EDM as a specific subfield of data mining. Four applications are grouped under “Student modelling”, six under “Decision Support Systems” and the last three are

Figure 2.1: Taxonomy of Applications in EDM



presented as “Other” because they differ from the other applications.

In the remainder of this section we will describe and illustrate these applications with the help of research examples related to each category of applications for more clarity.

2.3 Student modelling

Student modelling is a process devoted to representing cognitive aspects of student activities, such as analyzing the student’s performance or behaviour, isolating underlying misconceptions, representing students’ goals and plans, identifying prior and acquired knowledge, maintaining an episodic memory and describing personality characteristics [49, 12].

We have used this definition as a guide for categorizing some of the applications in EDM. All applications in this category present a model that describes

students attempting to reach their objective. Based on the literature review in 2013 [12], there are different characteristics in student modelling, namely, knowledge and skills (1), errors and misconceptions (2), learning styles and preferences (3), affective and cognitive factors (4) and meta-cognitive factors (5). Modelling student activities and behaviour can be used for predicting values representing students (the characteristics above) or discovering structures that describe students. As a result, there are two sub-categories in student modelling: prediction and structure discovery. In prediction, we usually know a specific attribute (representing characteristics) that we wish to predict and in structure discovery, we may not know the specific attribute or it may be only defined as a structure, instead of a single property. It is also important to mention that there might not be a clear line between these two sub-categories in all cases; but as there are enough differences in the objective of these two groups, it seems preferable to distinguish them.

2.3.1 Predicting Student Performance, Achievement of Learning Outcomes or Characteristics

In this set of applications/tasks, the objective is to estimate a value or variable describing students. This value can indicate students' performance, achievement of learning outcomes or characteristic. Most of the existing publications are focused on prediction of students' academic performance, but studies also look into characteristic such as collaboration with other students. The most widely used methods for predicting student performance and characteristics are regression and classification but other techniques have also been used such as clustering and feature selection. Miller et al. have used Lasso feature selection for identifying student characteristics which impact the learning. They compare the DM method with other models and examine if the selected features can be used for predicting student performance [33]. Zimmermann et al. introduced a model-based approach to predict graduate-level performance using indicators of undergraduate-level performance. Feature selection and prediction techniques have been used in this study [55]. Galyardt and Goldin have used the recent student usage data in order to improve the prediction

accuracy of the system in ITS [17]. Research by Waters et al. identifies the collaboration of students in online courses using Bayesian classification [54]. Sabourin et al. develop a model that identifies engagement of students based on their off-task behaviour in educational software. It investigates whether or not off-task behaviour can be a self-regulator of emotions [47]. Cocea and Weibelzahl estimate the motivational level of learners using decision trees [13].

2.3.2 Detecting Undesirable Student behaviours

This set of applications/tasks are similar to the prediction of student performance and characteristics, but in this category, the focus is mainly on detecting undesirable student behaviour, such as low motivation, erroneous actions, cheating, dropping out, academic failure, etc. The main data mining methods used in this category of applications are classification and clustering but other techniques are also applicable, such as feature selection and outlier detection. An example of this group of applications is the research done by Bravo and Ortigosa in which they propose an approach for detecting potential symptoms of low performance in e-learning using production rules [9]. In another study, Dekker et al. have used a decision tree classifier to predict student drop in an electrical engineering program [14]. Lykourentzou et al. used multiple machine learning techniques such as Support Vector Models and neural networks for students drop out prediction [29].

2.3.3 Profiling and Grouping Students

As the title of this category implies, the objective in this set of applications/tasks is to profile students based on different variables such as characteristics and knowledge, or using the information to group students for various purposes. Grouping students can be done based on various properties of profile information. This task is often different from clustering similar students with each other, as the purpose is to group students so as to complement each other. Also, when clustering students, one is looking for the greatest dissimilarity between clusters, but this may not be the case in grouping tasks. When using a grouping task for forming teams in a course project, one prefers to

have groups that are similar, while comprising dissimilar students that can complement each other. In a way similar to other categories of applications, different data mining methods can be used for these tasks, such as feature selection and clustering. As examples of this category, Azarnoush et al. proposed a method for learner segmentation using a dissimilarity measure based on a random forest [5]; Kinnebrew et al. used sequence mining techniques to identify learning behaviour patterns differentiating distinct groups of students [26]. Harley et al. studied the task of clustering and profiling of learners based on their interactions with an intelligent tutoring system [21].

2.3.4 Social Network Analysis

In this category of applications, the purpose is to obtain a model of students in the form of a graph, showing different possible relationships among them. In other applications of modelling, the focus is mostly on individuals, but in social network analysis (SNA), the focus is on the relationships between individuals. As an example, collaboration is a property assigned to the relationship between individuals, and to study it, one must model the relationships as well as the individuals. Rallo et al. used data mining and social network analysis to model the dynamics and structure of educative online communities [39]. Reffay and Chanier used social network analysis to measure cohesion in collaborative distance learning environments [40]. Reyes and Tchounikine studied structural properties of learning groups based on a relational perspective using social network analysis techniques [42].

2.4 Decision Support Systems

The other major group of applications/tasks in EDM is the decision support systems. Applications devoted to this category enhance the process of learning by helping stakeholders make decisions. Examples of this category are: providing feedback, creating alerts, planning, generating recommendations and enhancing the courseware. The target of these decision support systems is mostly the instructor, but it can also be the student, administrators or re-

searchers.

2.4.1 Providing Reports

Data analysis and visualization can be used as one part of many other applications, but it can also be an application by itself in educational environments by providing useful information to educators and administrators to help them with decision making. As a result, the purpose of this category of applications is to find and highlight the information related to course activities which may be of use to educators and administrators and provide them with feedback. The results of most of the applications grouped in “Student modelling” can be used for creating reports. Examples of this are: providing feedback on student performance or characteristics, describing the connections and collaborations through social network analysis and creating reports from the profile information extracted with the help of profiling methods. An example of this category of applications is the research done by Romero et al. in which they used association rule mining to provide feedback to instructors from the multiple-choice quiz data [45].

2.4.2 Creating Alerts for Stakeholders

This category of applications is similar to applications in the student modelling category. Usually, the purpose is to predict student characteristics and detect unwanted behaviour, but it mainly serves as an online tool for informing stakeholders or creating alerts in real time. Examples of situations in which alerts may be needed are in cases of low motivation, misuse, cheating, etc. An example of study in this category is the research of Knowles, which introduces a dropout early warning system using statistical models and regression [27]. In another study, Macfadyen and Dawson have developed an early warning system for educators using performance prediction [31].

2.4.3 Planning and Scheduling

The objective of this category of applications is to help stakeholders with the task of planning and scheduling. It can help educators and administrators with planning future courses or resource allocation, assisting in the admission and counselling processes or any other tasks involved in planning and scheduling [46]. It can also help students with course enrollment planning, in which case it has some common ground with recommender systems. In research with the objective of planning and scheduling, various methods have been used, such as discovery with models, cluster analysis and classification. Hsia et al. enhanced course planning by establishing the probability of enrollees completing their courses based on the student's preferences and profession [23]. The research of Delavari et al. discovers explicit knowledge that can be useful for decision making processes as well as proposing analytical guidelines for higher education institutions [15]. Huang et al. used cluster analysis, decision tree algorithm and neural networks for planning courses [24].

2.4.4 Creating Courseware

Courseware is known as educational software providing content, videos, tests and other learning materials. In this category of applications, the objective is to help the educator create or development course material automatically using student usage information. An example of this category can be found in the research of García et al., in which they propose a system for developing, improving and maintaining web-based courses using association rule mining and collaborative filtering [18].

2.4.5 Developing Concept Maps

Concept maps are “graphical tools for organizing and representing knowledge” [35]. In this category of applications, the objective is to develop concept maps of various aspects to help educators define the process of education. In other words, concept maps provide domain models to educators. They can help with mapping different concepts to each other (i.e. ascertaining relation-

ships). Examples of concept maps are: hierarchy of topics in course material, relationships of skills and test items, correlation of test items and knowledge components, etc. The research of Agrawal et al. presents a study navigator for studying electronic textbooks by creating a reference for concepts which students are reading [2]. As a further example, Lee et al. used an Apriori algorithm to develop an automatically constructed concept map of learning, provided to educators [28].

2.4.6 Generating Recommendation

In most of the surveys and books on EDM applications, generating recommendation is presented mainly as making recommendations to students. But recommendations can be targeted to any stakeholders. Examples of this category of applications are: course recommendations to students or test item recommendations to educators. The most common methods in recommender systems are collaborative filtering, content-based methods, association-rule based algorithms and hybrid approaches also used in EDM. Another method of generating recommendations is using discovery with models. For example, Vialardi et al. used a performance predictor model for generating recommendations [52]. The predictor model predicts the success of each student in each course and will recommend courses which the student is most likely to be successful in. In another study, O'Mahony and Smyth develop a course recommender system using collaborative filtering [36].

2.5 Other Applications

2.5.1 Adaptive Systems

This category of applications is related to the use of intelligent systems in computer-based learning, in which we need the system to adapt to the user's behaviour (this also referred to as "personalization"). This application is important, because in many online learning systems, numerous learners with different needs are involved with the system. And, as the number of participants grows, it becomes harder to meet the specific needs of all learners.

Adaptive systems can help us meet the needs of every individual learner. This adaptation can take on the form of adapting the course material, instruction pace, providing hints, ordering and generating tests, etc. As an example, the research of Alaofi et al. explores the personalization of a digital library using the student's profile information in order to improve search results [3]. In another example, Tang et al. propose a method for personalizing courseware construction using data mining techniques for distance learning environments [50].

2.5.2 Evaluation

Evaluation is one of the aspects of the educational environment which may not always be intuitive, especially in computer-based learning environments. As an example, evaluation in ill-defined domains has been researched in intelligent tutoring systems and is known as a domain that lacks a definitive solution (or the solution is dependent on the problem's conception) [30]. As a result, the evaluation in these domains is challenging. In this category of EDM applications, the objective is to provide an evaluator to help the educators. This can be done in exploratory learning environments and computer-based courses. An example of these applications is the research of Mallavarapu et al., which proposes a computational method to measure and track students' spatial reasoning in an open-ended simulation [32]. As another example, Hao et al. proposed a new method for scoring a game/scenario-based task using distance function [20].

2.5.3 Scientific Inquiry

Similar to other fields of study, theories and hypotheses on the process of learning and possible improvements are used in education. One use of educational data mining can be testing or even developing theories based on the various records in big datasets. This category of applications mostly targets the researchers as the end user, but any of the developed or tested theories can be used in other applications targeting other stakeholders later.

2.6 Literature Review

Data mining has been used for making discoveries in educational environments for the last few decades. In the last decade, the availability of online datasets and more uses of online learning systems have garnered more attention to this field. The Journal of Educational Data Mining (JEDM), created in 2009, has brought researchers of this field together. The publication of two books in recent years shows the growth of interest in this field, *Handbook of Educational Data Mining* published in 2010 and *Educational Data Mining: Applications and Trends* in 2013. There have also been multiple survey articles published about EDM so far. In this paper, we have attempted to integrate the ideas put forward in each of the surveys and books. In this section, we summarize the surveys and books related to applications of EDM, and the problems which have been addressed in each.

In the first survey published in JEDM and written by Baker and Yacef [6], four areas of application have been mentioned, i.e. improving student models (1), improving domain models (2), studying the pedagogical support provided by learning software (3) and scientific research into learning and learners (4). Student modelling in general or, as stated by Baker and Yacef, improving student models is one of most cited domains of research in EDM. We have unfolded this group into more detailed categories in the previous section. Improving domain models based on the objective of application (which is our focus) can be known as part of decision support systems such as developing concept maps or proving reports. Also, studying pedagogical support and scientific research can be summarized as scientific inquiry.

In another survey about EDM applications written by Romero and Ventura [46] and published in 2010, 11 categories of application were suggested based on 300 research studies completed before 2010. This survey has been extremely useful as a reference for this paper, as it provides many examples for each of the introduced categories as well as methods and techniques used in them. The categories of applications introduced in this survey are:

- Analysis and Visualization of Data

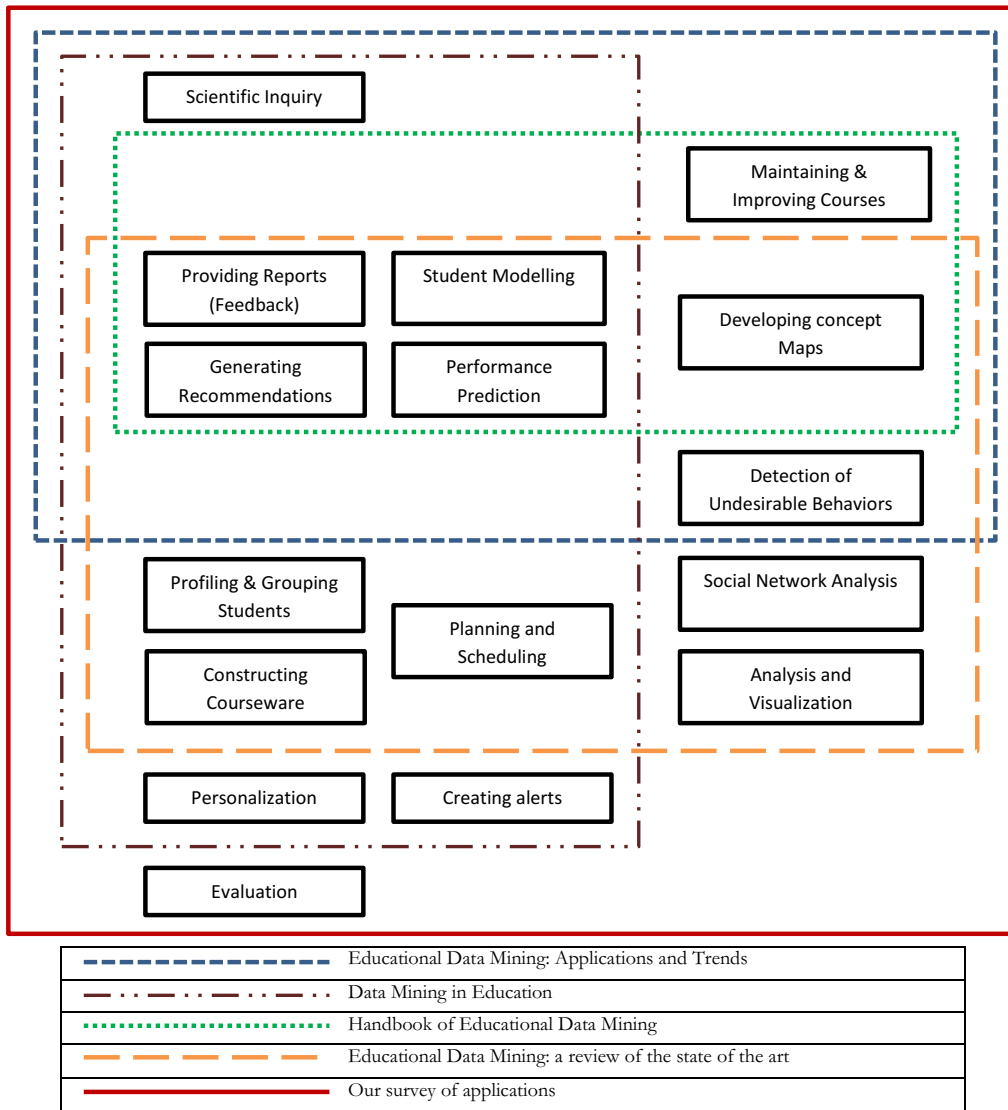
- Providing Feedback for Supporting Instructors
- Recommendations for Students
- Predicting Student's Performance
- Student modelling
- Detecting Undesirable Student behaviours
- Grouping Students
- Social Network Analysis
- Developing Concept Maps
- Constructing Courseware
- Planning and Scheduling

These applications are all mentioned in our list of applications with a few changes and additions. A part of the end objective analysis and data visualization is to provide reports (or feedback, as mentioned above). Moreover, providing reports (or feedback) is not simply limited to supporting instructors; it can also target students and administrators.

In another survey [43] published in 2013 by the same authors, the same list, with a few changes, is introduced. The main topic presented in this survey is the idea of categorizing applications and tasks based on the end user. This is potentially helpful in developing a better understanding of the problems related to EDM. The end users proposed by Romero and Ventura are learners, educators, administrators and researchers. We have used the same classification in our report. We have also specified the possible targeted end users of each application based on the same classification.

The Handbook of Educational Data Mining [44] was published in 2010, and discusses some of the research done in the field of EDM. The focus of this book is not to categorize EDM applications, although it does mention a few of them, namely communicating to stakeholders, maintaining and improving courses,

Figure 2.2: Comparison of Reported applications in EDM



generating recommendation, predicting student grades and learning outcomes, student modelling and domain structure analysis. These applications have been discussed in more details in the previous works of the authors (Romero, Ventura, Pechenizkiy and Baker), as mentioned earlier.

In 2013, the book *Educational Data Mining: Applications and Trends* [37] was published. It includes an introduction to the field which discusses different applications, methods and datasets. In the second, third and fourth part of the book, there are some case studies categorized in three groups (student modelling, assessment and trends). The general goals of EDM are introduced in the first part of the book to show the wide range of applications. The following topics are discussed.

- Student modelling
- Predicting student performance and learning outcomes
- Generating recommendation
- Analyzing learner’s behaviour
- Communicating to stakeholders
- Domain structure analysis
- Maintaining and improving courses
- Studying the effects of pedagogical support that can be provided by learning software
- Advancing scientific knowledge about learning and learners through building, discovering or improving models of the student.

This list shares many topics with previous works, but the case studies in the third part of the book entitled “Assessment” add something new. These studies cover evaluation in cases in which evaluation is not an easy task. Improving evaluation as a possible application of EDM had not been addressed before.

That was one important reason for us to specify a category of applications under “Evaluation”, as discussed earlier.

Finally in one of the most recent surveys on EDM in higher education [22] published in 2016, four categories of applications have been counted, namely course management systems, student behaviours, decision support system in higher education and student retention and attrition. This list of applications introduces a number of EDM tasks, but its treatment is not as clear, nor does it cover wide range of applications. For example, student retention can be considered as a part of student behaviour and there are many different topics in course management systems which may result in different applications.

To conclude, we would like to compare the different categories of applications in various surveys and books with each other as illustrated in Figure 2.2. In this figure, we show what applications each of the cited references cover. Some of the labels for categories of applications may vary, but in the end, they address the same applications using a different terminology. For example, we have equated domain structure analysis, cited in a few sources [37, 44] with constructing the concept maps. The definitions provided for these concepts are not exactly the same; however, based on the examples provided by the references, we concluded that they represent the same group of applications. This figure shows the topics that have been mentioned more often than others or the topics for which there has been less focus as on others.

2.7 Brief Summary

In this study, we reviewed the existing surveys and books about EDM and integrated the tasks introduced in each of them. We studied the examples provided in the surveys and books, as well as the publications of recent years in the *Journal of Educational Data Mining* to see if our proposed list covered all the research results. We grouped similar applications into categories and sub-categories in order to propose a taxonomy of tasks in EDM. This classification of applications/tasks is based on the end objectives. We chose some representative examples for each category. These examples can help us un-

derstand the categories better. However, they do not cover all the possible tasks in each category. We compared our proposed list of applications with the existing publications. Our list of applications is more exhaustive in terms of EDM topics compared to previous surveys and books, and proposes a novel and better suited categorization in a growing field. With the growth of using computers and availability of data, we believe the uses of EDM will also grow, leading to yet new applications.

Chapter 3

Graduating Attributes: Definition and Implementation

In exploring and setting up the curriculum and the programs of study and course objectives for their students, administrators in of higher education strive to identify and achieve the intended learning outcomes they have set. In order to describe and assess learning outcomes we would need descriptive attributes. These attributes should be able to describe the values and qualities which students will develop during their years of study. In 2011, University of Alberta's provost office published a report [1] describing the graduating attributes, the needs of an implementation and suggestions for its implementation. In this chapter we will describe the competency based education and graduating attributes. Then we will describe our implementation of Graduate Attribute Assessment Tool (GAAT) needed for assessing Graduating Attributes and collecting data.

3.1 Competency Based Education

Competence is known as “the ability to do something successfully or efficiently” (Oxford dictionary). It encompasses the knowledge, skills and abilities (KSA) which students learned during the years of study, and we expect individuals should develop to be helpful in their social life. In competency-based education (CBE) we try to master the learning outcomes as well as introducing strategies which provide flexibility to the process of learning [51].

This was introduced as complementary to our traditional educational system.

For clarity, we can look into some differences of competency based models and traditional models in aspects of pedagogy, adaptation and completion. In traditional models focus of teaching is on transferring information to students, but in competency based model those information and concepts are the means for building competencies. In terms of adaptation, students were needed to adapt to the pace of classroom. On the other hand in the CBE model the focus is on personalized learning. The other way CBE differentiates itself with traditional models is how it judges completion of learning process. In the traditional model the final grade will determine completion of course at the end of the term, but in CBE model, evaluation can be ongoing and based on qualitative assessment as well as final grades.

3.1.1 Design Criteria of Competencies

The question to answer after understanding the CBE is that what properties our design of competencies should have? Burke in a paper in 1974 [11] have counted a set of criteria for describing and assessing competency. That long list is describing features our targeted competencies should have, as well as the properties of educational systems in CBE. Some of those can be mentioned as:

- Competencies are based on an analysis of the professional role(s) and/or a theoretical formulation of professional responsibilities.
- Competencies describe outcomes expected from the performance of professionally related functions, or those knowledge, skills, and attitudes thought to be essential to the performance of those functions.
- Competencies are treated as tentative predictors of professional effectiveness, and are subjected to continual validation procedures.
- Competencies are specified and made public prior to instruction.
- Learners completing the CBE program demonstrate a wide range of competency profiles.

- Learner progress is determined by demonstrated competence.
- The extent of learner's progress is made known to him/her throughout the program.
- Competency measures are specific, realistic and sensitive to nuance.

CBE promises better outcomes and personalized learning to students; however, it is abstract and also costly for most of the institutions to change completely to CBE and meet all the criteria needed in CBE. That was the reason why some institutions moved toward the path of adaptation to CBE instead of a whole reform. As a result, we would need a list of essential elements of a CBE program. The 5 following statements describe the most important issues that can be used as a summary of competencies criteria. These essential elements were described by the early influential work of [16] which is listed below.

1. Competencies are role derived, specified in behavioural terms and made public.
2. Assessment criteria are competency based, specify mastery levels and are made public.
3. Assessment requires performance as prime evidence but takes knowledge into account.
4. Individual student progress rate depends on demonstrated competency.
5. The instructional program facilitates development and evaluation of specific competencies.

Looking over these essential elements there are two topics which are outstanding in CBE. First the way we describe competencies to the stakeholders of educational system and second, the way we assess competencies.

3.1.2 Assessment

“Assessment is the systematic collection, review, and use of information about educational programs undertaken for the purpose of improving student learning and development.” [7] As a result, assessment is an important aspect in the context of education. Based on the usage, there are different types of assessments, following different purposes.

- Formative assessment: Assessment used to monitor student progress during instruction.
- Diagnostic assessment: Assessment focused on the common sources of error encountered by student, so that learning difficulties can be pinpointed and remedied.
- Summative assessment: Assessment at the end of instruction for the purpose of certifying mastery or assigning grades.

Assessment can also aid learning by (1) motivating students, (2) increasing the retention and transfer of learning, (3) helping students to understand themselves better (self-assessment) and (4) evaluating the effectiveness of instructional process [19].

The assessment we are using in this research can be counted as the formative assessment. As introduced in transformative assessment book [38], “formative assessment is a planned process in which elicited evidence of students’ status is used by teachers to adjust their ongoing instructional procedures or by students to adjust their current learning tactics”. [38]

First we should understand that the assessment we are using is a planned process. In other words, we should know at which stages we are going to use this assessment and it will be available to students during the whole process of learning. In our case we will use the assessment for all years of student’s study. Our expectation is that the assessment will be done at the end of all semesters as well as the beginning of the first semester. Students are allowed to do the assessment more frequently and we will store their responses in our database for further analysis.

The adjustment to students' or instructor's plan and behaviour shall be based on the evidence provided by the assessment. The evidence in case of our model is the aggregation of students' self-assessment in terms of competencies and graduating attributes. The adjustment can be done in different ways. The students for example can get an idea of which competencies they need to focus on more and based on that they can plan their courses. On the other hand for the instructors the comparison of their expectation with the aggregation responses of students can provide some ideas, for adjusting the course planning.

3.1.3 Self-assessment

The definition of self-assessment provided by Andrade and Du [4] can be helpful in understanding of our model.

“Self-assessment is a process of formative assessment during which students reflect on and evaluate the quality of their work and their learning, judge the degree to which they reflect explicitly stated goals or criteria, identify strengths and weaknesses in their work, and revise accordingly.” [4]

As mentioned earlier in Section 3.1.2 as well as in the definition above we are considering the self-assessment as formative assessment. This self-assessment is both for students and instructors. In case of students, it informs us of the level of mastery for each Graduate Attribute (GA) acquired by the students. In case of instructors, it is more like a course assessment in terms of what instructor expects as the result of each specific course. As a result, the assessment informs us of the plan instructor has for each course.

Considering the rest of self-assessment definition, each student may be able to judge their degree of strengths and weaknesses. In the context of GAs, self-assessment may be hard as we are mostly talking about abstract cognitive attributes. To help with this issue we have provided a definition of each value for each sub-attribute, explaining what we mean by each value for each sub-attribute. As a result of these explanations as well as levels of acquisition, users share similar understanding of values. This can help us with gathering more reliable data from the assessments. We will discuss the details of self-

assessment tool more in Section 3.3.

3.2 Graduating Attributes

Graduating attributes try to describe the outcomes from a more general perspective compared to competencies. For a better understanding we can look at Graduating attributes (GAs) as defined by Bowden et al.

“Graduating attributes are the qualities, skills and understandings a university community agrees its students should develop during their time with the institution. These attributes include but go beyond the disciplinary expertise or technical knowledge that has traditionally formed the core of most university courses. They are qualities that also prepare graduates as agents of social good in an unknown future.”[8]

Based on this definition and the GAs used in this study, we can define a list of properties graduating attributes have.

- Can describe the skills, qualities and understandings students are going to develop.
- Administrators, instructors and students should agree that the list of GAs is an outcome of education.
- There should be assessment based on the specified list of GAs.
- The list of GAs should include qualities which prepare students for future life in society or work environment.

This project uses the list of GAs suggested by University of Alberta provost office [1]. Later on, the research done by Ipperciel and ElAtia [25] from University of Alberta proposed an assessment model for the GAs. The focus of this research was on presenting a criteria-based model for assessing GAs. In this model we have seven attributes, each having four sub attributes. These attributes are, Ethical responsibility(1), Scholarship(2), Critical thinking(3), Communication(4), Collaboration(5), Creativity(6) and Confidence(7). Below, we will discuss the properties this model provides.

3.2.1 Scales with intuitive labels (GA Levels of Acquisition)

The proposed approach consists of a scale with intuitive labels. Levels 1-2 correspond to levels of pre-acquisition. Levels 4-5 designate levels of excellence that may go beyond what is expected in a university setting and may not be reached by all students. Labels proposed are emergent (level 1), basic (level 2), adequate (level 3), superior (level 4) and exceptional (level 5). At the emergent level, GA acquisition refers to the awareness of individual (or atomic) elements that are needed to perform a specific task, such as the cognizance of facts, ideas or rules. At the basic level, GA acquisition involves manipulation and combination of the basic individual elements (facts, ideas or rules) in a coherent (molecular) whole aimed at performing a specific task. At the adequate level, the minimum standard/norm for performing a specific task has been met and a GA is deemed functional in the academic context. At the superior level, a GA is acquired to the extent that it allows for new applications in, and generalizations to, unforeseen contexts. Finally, at the exceptional level, there is a consistency and spontaneity in the capacity to generalize GA application and adapt to new situations, including outside the academic environment. [25]

3.2.2 Generic attributes Interpreted by a series of sub-attributes:

As discussed in the previous study [25] and the provost office of University of Alberta [1], institutions tend to define a short list of graduating attributes which is more generic. As a result, there is the need of defining a series of sub-attributes for each of generic GAs. Given the many disciplinary cultures within a large institution, and the different requirements of specific programs, there is often a need to define a specific interpretation of the sub-attributes peculiar to a faculty or a unit. However attributes and sub-attributes are general concepts applicable to the whole university. The proposed model in this paper consists of seven generic attributes each specified in four sub-attributes presented in Table 3.1. This is the model we used in our research as well.

Attributes	Sub-attributes			
Ethical responsibility	Global citizenship	Community engagement	Social and environmental awareness	Professionalism
Scholarship	Knowledge breadth and depth	Interdisciplinary	Life-long learning	Investigation
Critical thinking	Analytic and synthetic reasoning	Interpretive proficiency	Intellectual curiosity	Information literacy
Communication	Writing skills	Oral skills	Visual communication	Multilingualism
Collaboration	Openness to diversity	Interpersonal skills	Adaptability and compromise	Individual contribution
Creativity	Imagination	Innovation	Divergent thinking	Artistic sensibility
Confidence	Leadership and empowerment	Independence	Initiative	Resilience

Table 3.1: Graduating Attributes and sub-attributes

3.2.3 Appropriate description of acquisition levels with regards to KSA GAs:

The GAs can be categorized in three categories of knowledge-type GA, skill-type GAs and attitude-type GAs. As different attributes and sub-attributes can refer to any of these categories, it is needed to have appropriate descriptions for each of the attributes and sub-attributes to indicate to which of the three categories they belong. In the model used in this project based on the previous work [25], descriptions have been provided for each of values which can be assigned to each of sub-attributes. This can be described in different levels of acquisition based on the type of attribute. As indicated in the paper this model needs to be implemented to study needs, interest and concerns of all stakeholders and propose useful changes to the model. That will be part of what we try to do in this study.

3.3 Online Assessment Tool

The next step in this study is to design and develop an online assessment tool considering the described Graduating Attributes for a an educational institution. In this online assessment tool we want to associate a value to each

of the GAs for each student and each course. Having this data alongside some profile data of the user can be a foundation for further research in educational data mining. Also, we need to know the possible challenges and consider them in our design. Related works shows us some important challenges about using GAs in application, listed below:

- Members of the institution do not share a common understanding of GAs.
- Stakeholders do not share the sense of necessity in use of GAs which can causes poor participation.

Previously we talked about using detailed descriptions for GAs and values assigned to them. These descriptions can help users in having a more common understanding of each GA. Also, starting the study with volunteer users can promise more engagement than a mandatory case study. This way we are trying to tackle the second mentioned challenge.

Our online assessment tool for graduating attributes will target three types of stakeholders. The first group of users are students who will assess themselves in terms of graduating attributes. The second group are educators who may use the tool to describe their courses in terms of Graduating Attributes. The other target of GAAT will be the administrators who will have access to the data gathered by the system for high level analysis and decision making.

We have implemented our tool as a Moodle plugin. Moodle is a free online learning management system. Moodle has been used in University of Alberta for managing courses, and it provides an online tool for managing course activities. Implementing our tool in Moodle has various advantages; (1) being linked to educational profile of users (students and instructors), (2) having access to the courses each user is enrolled in and (3) making it easier for users to use the plugin (they just need to use a single web application for all educational purposes). The first step in use of the tool is a questionnaire for each user, which can help us with gathering some profile information. Then they will have access to an assessment tool for assessing GAs. For students it is a

self-assessment in terms of GAs and for instructors it is a course assessment at the beginning of the semester. In the sub-sections below we will discuss each of these tools.

3.3.1 Online Questionnaire

The first step toward gathering useful data, for further analysis is profile information. As a result we have designed two questionnaires, one for students and one for educators. For designing such a questionnaire we first need to have an idea of the way we plan to use this data. Assuming that we have collected the response of students in terms of GAs (as we will discuss in later sections), there are a few questions which can be answered if we have some profile information. For clarification, in AppendixA there is a list of questions which can be answered if we collect respective information.

To answer these kinds of questions, we have designed a questionnaire for students and instructors to fill at the beginning of each semester. This questionnaire the same way as other parts of GAAT can be accessed at any time during the semester and can be updated later on.

3.3.2 Student's Graduate Attribute Assessment Tool (GAAT)

The first question in the design of this application is how we can assess students in terms of GAs. We have chose to design and implement the assessment of Graduating Attributes as a self-assessment tool. This means that students are going to assess themselves via our online tool in terms of GAs. One important aspect of the data in this project is the changes in GAs' values throughout the years of study for each student. To consider this aspect, the students will do the assessment at the beginning of the first semester and also at the end of each subsequent semester. This can help us with gathering information about the changes each student experiences through self-assessment during the years of their program.

Figure 3.1 shows a snapshot of student's GAAT. Each aspect of this tool will be discussed in more detail in the next sub-sections to show how we ended

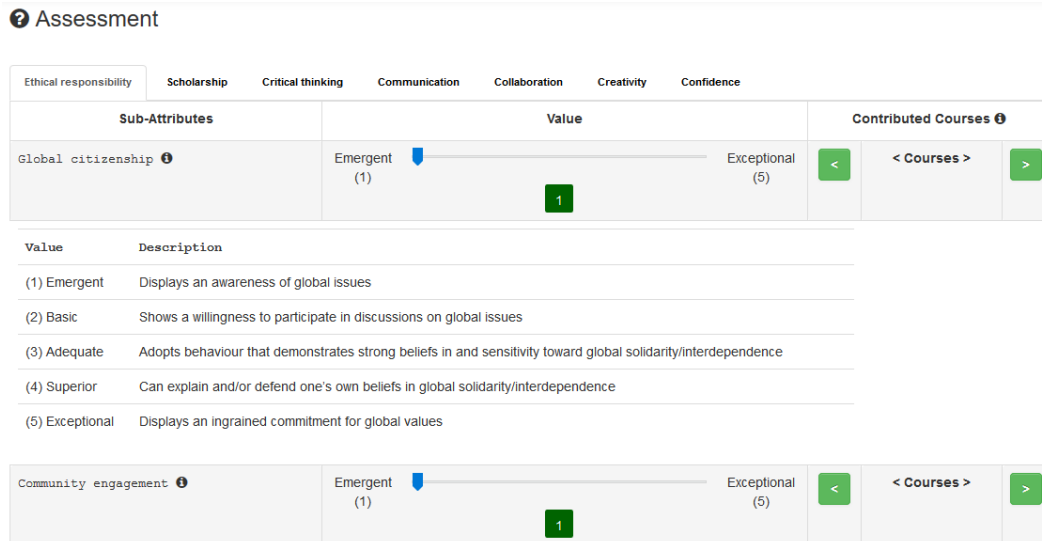


Figure 3.1: Student's Graduate Attribute Assessment Tool Snapshot

up with this design.

To describe students with GAs we need to have a value assigned to each of the sub-attributes. Also we need to show that each of these sub-attributes describes a part of the main attribute. For example the value assigned to "Global citizenship" describes a part of "Ethical responsibility". To satisfy this need we have grouped all sub-attributes of each attribute in a tab titled after the main attribute's name. This automatically shows the relations between sub-attributes. The same design has been used for the instructor's GAAT

In the students' GAAT we will ask users to assess their level of mastery by choosing a value between 1 and 5 for each of the sub-attributes. To show that these values are in a progressive order we chose to use a slider input. These values in order are titled, "Emergent", "Basic", "Adequate", "Superior" and "Exceptional" as described before. As mentioned before, there is a description for each value assigned to each sub-attribute, to help users have an understanding of what each value means for each sub-attributes (different values have different meanings based on the sub-attribute). As a result the description for values will be shown to the user respectively.

Aside from the values for each sub-attribute, there should also be a way to make the connection between the assessment and the courses of the students.

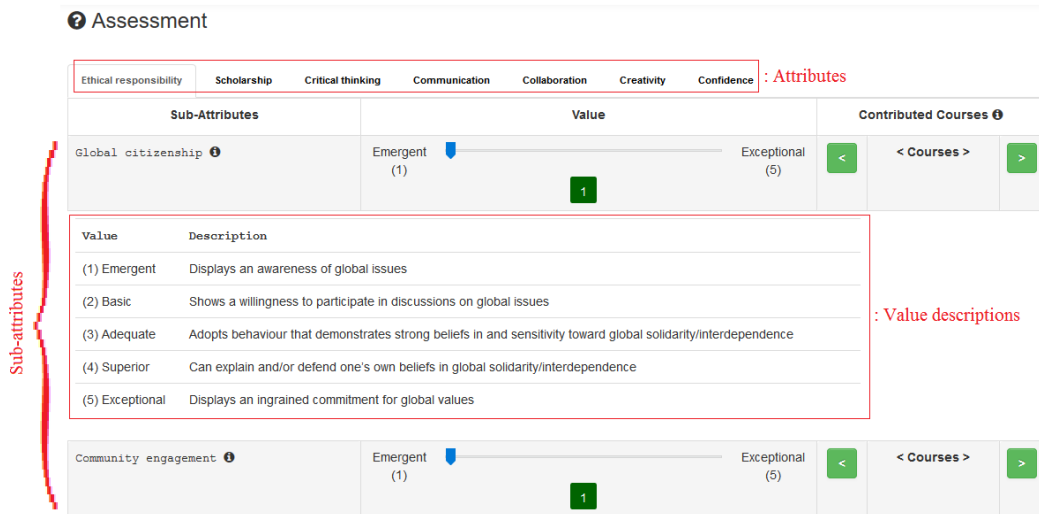


Figure 3.2: Student’s GAAT Snapshot showing attributes, sub-attributes and attribute descriptions

We expect that some of the improvements in GAs would be as a result of the courses. But the only way to make sure is to ask the users to tell us which courses had an impact on the improvement of the GAs. The first way to gather this information was to have an assessment for each of the courses. But there are two issues with this. Firstly it will be too time consuming for the users, since the time needed to fill out each sub-attribute assessment will be multiplied by the number of courses. Secondly, it may give students the impression of course assessment; however we are looking for a self-assessment in regard to GAs.

The other design which we use in the project is to have a column for each course in the GAAT. After assigning a value to a sub-attribute, users can choose all the courses which contributed to the assessed sub-attribute. Also, to gather more information we added two possible choices for each course, namely, major impact and minor impact. In order to gather mentioned information, we included a column named “Contributed Courses” in student’s assessment tool. As highlighted in Figure 3.3, students can go over all their enrolled courses of their current semester, and choose the impact of each course on the sub-attribute they are assessing.

In the first student assessment there will be no column as “Contributed

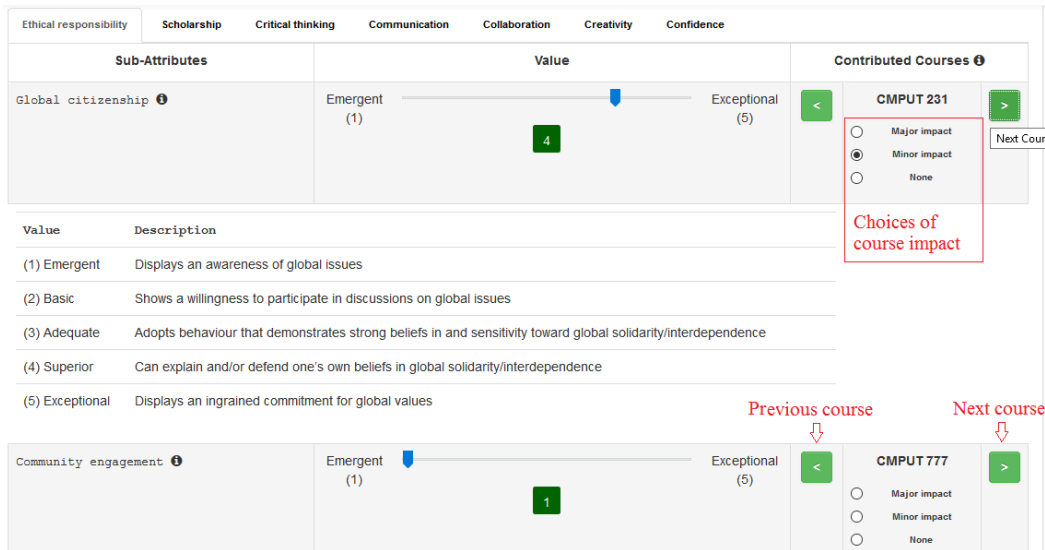


Figure 3.3: Student's GAAT Snapshot showing the contributed courses

Courses”, as we are just interested in the students’ status before they have been affected by their education. We will use this first assessment as a baseline for students’ improvement. After this assessment, all other assessments will be at the end of the semester and they include the “Contributed Courses”; so that we can see students’ self perceived progress as well as the courses’ influence on these changes.

The improvements in attributes of each student can be as a result of different activities. We divide them into two groups: the courses students enroll in and all others as extracurricular activities. The best scenario for gathering data was to ask about these information at each step of assessment (in each semester). In order to achieve this goal, we needed to first ask about student’s self-perception, and then observe improvements as a result of each course they had during the semester (asking for an assessment per course). Aside from being time consuming, this design needs a method to aggregate different assessment into one describing the student in that semester. Instead, in our model, we ask students to assess themselves and also mention which of their courses contributed to the sub-attribute they are assessing. This way, by comparing with previous assessment, we can figure out the sub-attributes which have been improved as well as the courses contributed to each of them.

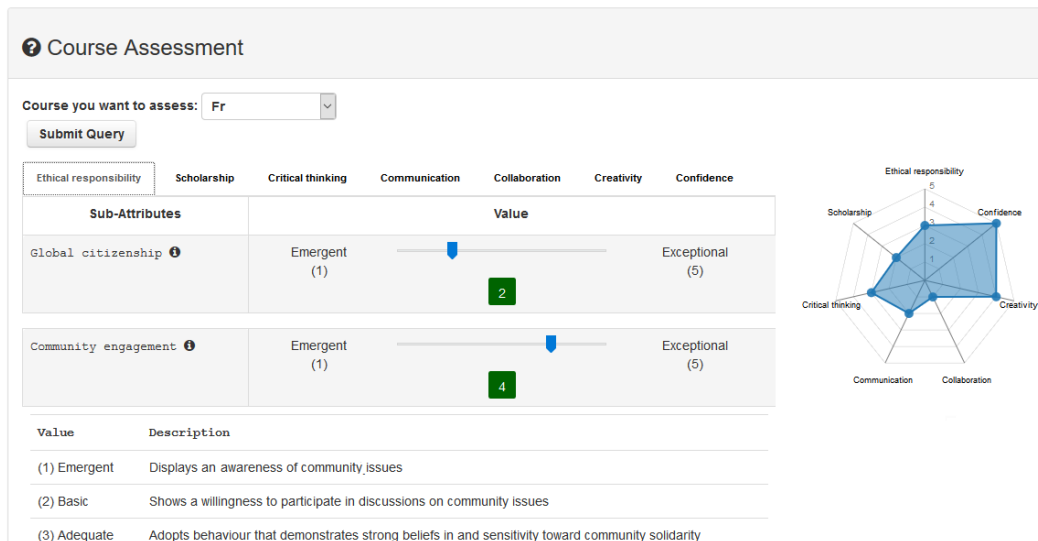


Figure 3.4: Instructors Graduate Attribute Assessment Tool Snapshot

3.3.3 Instructors Graduate Attribute Assessment Tool

The other data we are interested in is what the expectation of instructors is from the courses they present. In the instructors' GAAT we ask the educators to indicate which GAs and to what extent they are focusing on. This can help us to see if instructors' expectation is compatible with how the student's describe themselves. Also later on this course assessment can be used as profile information of courses. One of applications of this project can be course recommendation to students. To do this recommendation we need some values describing each course. Instructor's GAAT can help us with providing the descriptive values. A snapshot of the tool is shown in Figure 3.4.

The same as students' GAAT, instructors can assess their course, by assigning a value between 1 and 5 (emergent to exceptional) to each sub-attribute. The same descriptions for values will be used, to help instructors in their assessment. The graph shown in the right side of the GAAT will show the value assigned to each GA. This will be the average of sub-attributes for each GA. The purpose of this graph is to show the instructor a comparison between the GAs. Looking at this graph will give us insight about the course's focuses. Later on the same radar graph will be used in a report to instructors, comparing their expectation of the course with the cumulative responses of students

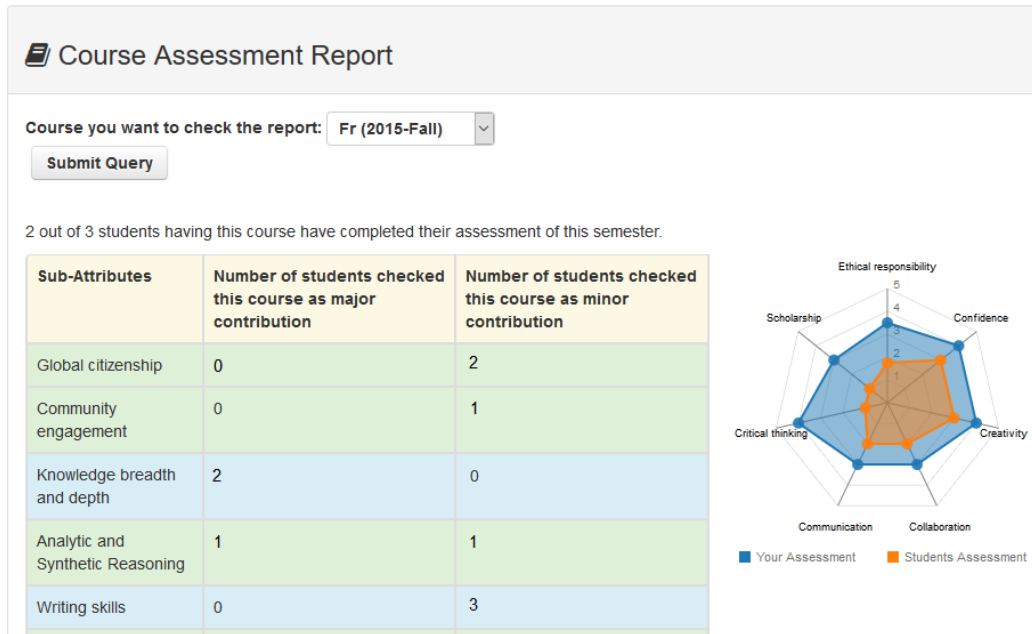


Figure 3.5: Course Report for Instructors

enrolled in that course.

3.3.4 Course Report for Instructors

Providing reports to instructors is useful in various ways. First it gives the instructors incentive to participate in this program, and second, will give them some feedback on the students enrolled in their classes. Also it can help them improve their course material respective to the needs of students. Figure 3.5 shows a snapshot of the report which will be automatically provided to the instructors via the Moodle plugin.

As mentioned before in student's GAAT, students will indicate the contribution of each course for each sub-attribute. The first part of this report shows the number of students indicated the course as having major or minor impact. This provides some feedback to the instructors about how the course has been useful for students. In the second part of the report, the course assessment of the instructor (in a radar graph) will be shown on top the average self-assessment of the students. This can help instructors to indicate which GAs needs to be more or less emphasized on.

3.4 Brief Summary

As described in this chapter, we designed a plugin for collecting data and providing some tools to the stakeholders. This design can be used as a platform for any more complex usage. As mentioned in Chapter 2, there are many possible applications in educational data mining, and some of them can be used in this project, in the direction of Graduating Attributes.

In the next chapter we will describe some of these possible applications in more detail.

Chapter 4

EDM Applications of Graduate Attribute Assessment Tool

4.1 Introduction

In Chapter 2 we looked over different tasks and applications in Educational Data Mining. And in Chapter 3 we described our design of the Graduate Attribute Assessment Tool (GAAT). Now we will look into the possible applications of GAAT including data mining applications. We need to identify applications for each of the stakeholders in our system. Useful applications for each of the stakeholders can help with the process of engaging them with the GAAT, by encouraging them in using the system. Also not having appropriate applications for each of the stakeholders may result in less accurate data as the users will not be encouraged to spend time on it. For example if there were no applications targeting the learners, they would not spend enough time doing the self-assessments and as a result the system would not be able to accomplish its purposes. As discussed in Chapter 2, there are four stakeholders we focused on namely learners, educators, administrators and researchers. While these stakeholders were discussed in previous chapters, there exists a hidden stakeholder namely the employers which is not directly involved with the system but can use it indirectly. Employers/Industry can use the GAAT indirectly by describing their job positions in terms of graduating attributes, and matching them with the existing students. One of the intentions of an implementation for GAAT has been enhancing the process of employing new

employees by helping both near graduate students and employers. With these five stakeholders in mind, we can think of many possible applications.

Starting from the educators, GAAT will help them with providing some feedback in the shape of reports about how the students have assessed each of the courses (as discussed in Chapter 3). The other possible application can be the grouping of students. This can be helpful for both learners and educators; when there is the need of forming groups. Grouping students is a task in which we need to match students in a way that each student's abilities complement group's abilities in order to have balanced groups. Graduating attributes and the record of students in the GAAT are the data we can use for accomplishing this task.

Targeting the administrators, we can provide more overarching feedback about how the courses in each department cover the graduating attributes and how successful they are in covering them. This information can help the administrators of the institution in designing new courses or maintaining the existing ones.

GAAT also provides a platform for the researchers working in the education field. Significant research exists about the use of competencies and graduating attributes which needs data for evaluation and development of further researches; GAAT can help with collecting these data.

For learners, we need to provide some useful applications to encourage them in using the GAAT. As a result one useful application can be a course recommendation system. In the GAAT we are recording the responses of different students over times, and based on the students responses over time we know that which of the courses can be more useful in developing each of the attributes. We can use these data by generating recommendations to students, using the previous experiences. This also encourages students to do a more accurate self-assessment as it can help with the more accurate recommendations.

Finally for the industries which are interested in hiring graduated students, there exist applications enhancing this. Employers can describe their job requirements using Graduating Attributes. By having these descriptions, we can

create a match making application working on providing the best people for each position. This can help both sides as it makes it easier for the employers to find the best fit for their positions, and it helps the graduated students to find the best match based on their abilities.

In this chapter we will provide a brief introduction to the two latest examples. The main focus of our research is generating course recommendation to learners and we will discuss this task in more details (algorithm and implementation) in the next chapter. But it is of importance that the implemented GAAT has other applications, which can be the main direction of future studies.

4.2 Matching Employees and Employers

First we need to define the application's purpose. One of the complications of today's educational environments is making the connection to the employing companies in regard to learning attributes. However grades can be used to represent the quality of learning the course materials in higher education institutions, there exists no general values to describe the attributes which students have developed in their years of study. If we could come up with such values, one of the useful applications could be a faster and more convenient way of matching employees and employers. For example once we have representative values for graduating attributes, employing companies can describe their requirements in terms of graduating attributes to the institutions and by having these list we can design a matching application. The matching can follow various criteria, such as matching each student to the best fitted positions (maximizing the similarity), matching each student to the position s/he have the best chance of getting it and combinations of them. In the next section we will briefly describe a possible design of this application.

4.2.1 Summary of Design

One way to describe students' status with graduating attributes is to represent each student with a vector of values for each of the existing graduate sub-




Figure 4.1: Example of matching students to job positions

attributes. The same can be done with the positions provided by employers. The requirements of each position can be represented with a vector of required value for each of the sub-attributes.

The next step would be defining a distance function to measure similarity of two vectors. One possible way is using cosine similarity measure. As a result it is easy to find the most similar position to each student and using it as the recommendation (nearest neighbour method). It is of importance that if a position is the closest position to a student, that student may not be the closest student to his closest position. For example in figure 4.1 the most similar position to S_1 is P_2 (similarity of 7), but there exists a better fit for position P_2 with similarity of 10; as a result P_1 might be a better match for S_1 . In order to consider this possibility the other option can be the reverse nearest neighbour method in which the algorithm will look for the student closest to each position.

4.3 Course Recommender System

The other application of EDM for GAAT is a course recommender system. In this application the purpose is to use recorded usage date of students and instructors in order to recommend future courses. There are various useful data to use in this regard, the profile of each course created by the instructors, the self-assessments of students and the evaluation of students about the impact of each course on their improvement (the tools introduced in Chapter 3).

The mostly used recommendation systems are collaborative filtering and content based methods. In short collaborative filtering uses the similarity of users' behaviour or items (in our case courses) for creating recommendations

and content based methods use the profile of items and try to choose the best of them based on the taste of the user. In the next chapter we will describe these methods in more details. In our application we can use both of these methods (separately or together). The main data needed for collaborative filtering method is to have a rating to each of the items. We can use the impact of courses reported by student in place of the rating. In the user based collaborative filtering the system would be able to find the similar students, and then for generating a recommendation uses the high rated items of similar users. Or in the item based collaborative filtering, system will find the similar items to a previously high rated item and will recommend those. In order to use content based methods in our system, the assessment done by the instructors will be used. In the simple form if we want to recommend a course to a student in order to improve his collaboration attributes, we will recommend the course which have be given a bigger value to collaboration by its instructor.

The focus of the next chapter will be describing the algorithm capable of recommending courses based on graduating attributes and evaluating its efficiency.

Chapter 5

Course Recommender System

5.1 Introduction

As we discussed in Chapter 4, one of possible applications of gathering data related to Graduating Attributes (GAs), via our Graduate Attribute Assessment Tool (GAAT) is recommending courses to students. On the one hand this can help students with choosing courses and on the other hand it gives students a reason to engage with the process of gathering data.

There exist two main methods used for development of recommender systems, collaborative filtering and content based methods. Collaborative filtering uses the historical preferences of users for making recommendations. There exist multiple ways of using collaborative filtering. As an example, in user-based collaborative filtering, we calculate a similarity between users using their historical preferences. Then for recommending new items we look into the items which users similar to our target user prefer. In the context of our application, the items will be courses. We can look into the ratings given to courses to calculate similarities of students and using them for making recommendations. We will discuss the algorithm we have used for our course recommender system in Section 5.3.

The first challenge with the implementation of this application is that for developing an appropriate course recommender we need to have some data. However, we currently lack this data. This led us to the idea of generating synthetic data for training and testing our algorithm. In Section 5.2, we will describe the properties of our expected data and the model on which we

generated our synthetic data.

5.2 Data

5.2.1 Structure of Data

We have discussed in the previous sections the modules which will gather the data. Questionnaires designed for students and instructors (students' survey and instructors' survey), students GAAT and course assessment for instructors. Apart from these four collections of data, in the online management tool we have some profile data describing users (both students and instructors) and courses. The relations between these collections have been described in Figure 5.1.

For each student we can have multiple responses to the questionnaire (students' survey) and also for each student we can have multiple responses to their self-assessment. Each of the assessments done by the student can relate to multiple courses (courses contributing to the improvement of GAs). The same relational structure exists for the instructors.

Depending on the scenario and the gathered data, we can describe the needed data entities for the task of course recommendation. First, we need to have a list of Graduating Attributes and a range of values to describe them. Then, there would be a list of courses available for the students. Each of the courses needs to be described in terms of Graduating Attributes, meaning that they will improve each of the competencies by how much from the instructors' point of view. Each student will also be described by the list of values assigned to the Graduating Attributes. In each semester each student would participate in some of the existing courses and those courses would improve the values of competencies. Based on how impactful the courses have been in increasing the value of each Graduate sub-attribute, the student will rate it. By having these data, our application will generate recommendations for courses in the next semesters.

For implementation and evaluation of a course recommender system we need to have the data. As we were not able to gather a sufficient amount

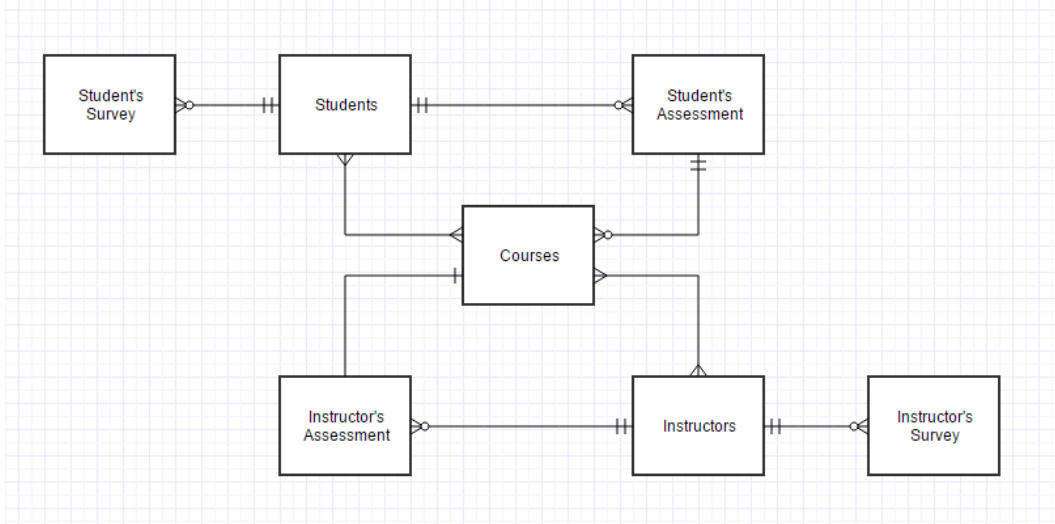


Figure 5.1: Relations of data entities

of data during our case study, we were left with the option of generating synthetic data. We will describe the process of the data generation in the following sections.

5.2.2 Data Generation

To generate realistic data for the use of our applications, we need to consider how the real data would look like. We have formulated the following list of assumptions and we have considered them in our data synthesis. The list of assumptions we have used is as follows:

- Students may start with different values for sub-attributes as a baseline (They may not be similar), so we can't start all from the zero value for all attributes. The random value assigned to each sub-attribute for each student is independent of other values.
- Personality of students is a factor in the assessment of competencies (some may report most of the courses as major impact and some the opposite). Students' capacity of learning each sub-attribute is also different from other sub-attributes. Students' personality is independent and constant over time.

- The probability of values assigned to Graduate sub-attributes for each course is not uniform. It is more likely to have lower values compared to higher values.
- In the updating of the values of competencies, all the courses taken in the semester have impact. We assume that the students spend equal time on each course, so we sum the impact of all courses and we divide it by the number of courses.
- The number of courses taken by a student in each semester is not constant. We have assumed that students will take randomly between 4 to 6 courses per semester.

For this study we have assumed having 100 students, 100 courses and 28 sub-attributes (likewise the GAAT). Each of the competencies can have a value between 1 and 5 (from low expertise to high expertise). Each course needs a list of values to describe the amount of focus on each of the sub-attributes (C_{kj}). This value has been assigned randomly with more focus on lower values. The value will be between 1 and 5, but the probability of values being between 1 and 2 is 36% and it decreases to 14.5% for values between 4 and 5. We have used such a value as we expect less courses with values of 4 and 5 compared to lower values. Figure 5.2 shows the different focus of 5 courses in the generated data. It shows that there exist less values in the range of 4 and 5 compare to the values in range of 1 and 3.

Also we have assigned a uniform random value between 1 and 3 to each sub-attribute of each student as the baseline of student. This way we are considering the different background of students assuming that no student will start having the highest value which is between 4 and 5. In order to create different personalities of students in our data, we have assigned some other values to each student which based on those we would create the rest of the data. Each student has a talent ratio for each of the sub-attributes, meaning that different students may have different learning potential in regard to each of the sub-attributes (T_{ij}). Also students may act differently in terms of rating

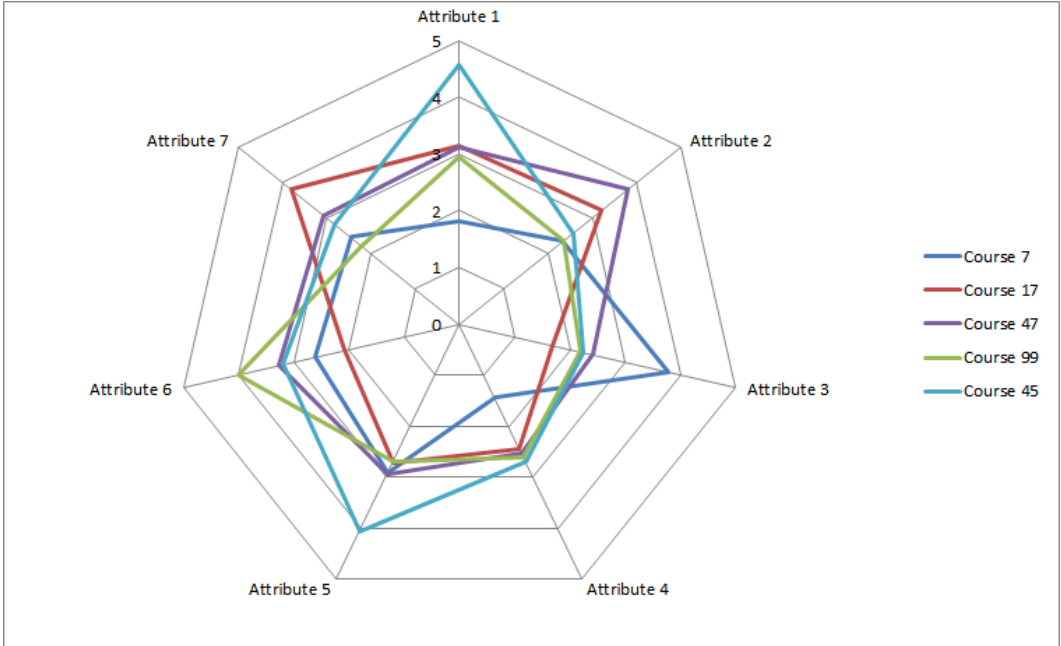


Figure 5.2: Values assigned to some example courses in the generated data

their courses. The rating in our synthesized data is rather 0, 1 or 2. Some of the students may rate small changes as 1 but some others may start giving the rating of 1 to bigger changes. As a result for each student we have assigned two limits between 0 and 1 to determine what are the limits of their rating ($0 < l_{i1} < l_{i2} < 1$). If the impact of each course on a competency is between 0 and 1, each user i will rate it 0 if it is less than l_{i1} , 1 if it is between l_{i1} and l_{i2} and 2 if it is more than l_{i2} . Using these values which represent the personality of each student we will create the rest of the data. We should also mention that the values used for representing the personality of students are independent of each other.

In each semester, students take 4 to 6 courses randomly from all the courses they have not taken before. The new value assigned to each sub-attribute for each student at the end of each semester is calculated with equation 5.1. As a result each of the courses will have some specific impact, and by comparing that impact with the l_{i1} and l_{i2} we will create the students ratings as well. Figure 5.3 shows the increase of the value assigned to a sub-attribute for different student over time. This can show that students start with different values and

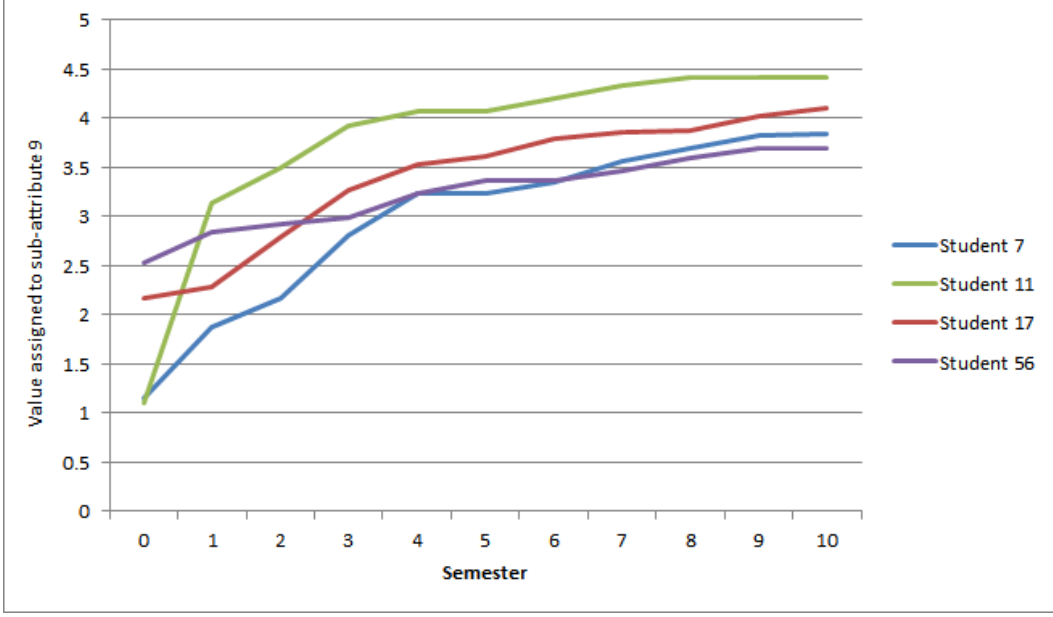


Figure 5.3: Values assigned to an example sub-attribute over time for some example students in the generated data

the value increases differently for each student over time. For example student 56 has started the first semester with one of the highest values 2.5 compared to other students, but has improved the value with a low pace and has the value of 3.7 by the last semester. On the other hand student 11 has started with value 1.2 but has the value of 4.4 by the last semester.

$$B_{ijt} = B_{ij,t-1} + \sum_{k=1}^K \frac{\max(0, C_{kj} - B_{ij,t-1})T_{ij}}{K} \quad (5.1)$$

- B_{ijt} represents the value assigned to sub-attribute j for student i in semester t
- C_{kj} represents the focus of course k on sub-attribute j
- T_{ij} represents the talent ratio of student i for sub-attribute j
- K represents the number of courses for student i

The increase over the values of sub-attributes for each student also results into lower rating in the latest semester as there would be fewer courses to increase those values. Figure below shows the ratio of different rates in each




Figure 5.4: Average ratio of ratings for all student over time

semester. Over time the average number of 0 ratings increases and at the same time the average number of 2 ratings decreases.

5.3 Algorithm Design

We will first describe the utility of a classic Collaborative Filtering (CF) method. By this term we refer to the algorithm that bases its predictions on neighbours of relevant users. The idea of CF methods lies on the fact that people who agreed in their evaluations for past items are likely to agree again for future items [41].

As for the format we use, each student is represented by a long vector of the courses taken and the assessments (ratings) in regard to the increase in his/her personal Graduating Attributes (GAs) achieved. Similarly, each course is represented by a baseline (ground truth) for the degree that student competences are supposed to increase after the course is taken. One variation in our method of formatting the problem with other applications of CF, is that each item is a pair of Graduate Attribute and a course. This way we can not only recommend a course to a student considering most increase in all GAs, but also

recommend a course only targeted on a specific GA. In a hypothetical example of $C = 11$ courses ($cid = 0, 1, \dots, 10$), $S = 5$ students ($sid = 0, 1, \dots, 4$) and $G = 4$ competences ($gid = 0, 1, 3$), the matrix representation of the problem would look like Table 5.1.

Table 5.1: Example of data representation: courses x competences x students

cid					sid=0				sid=1				sid=2				sid=3				sid=4			
	g0	g1	g2	g3	g0	g1	g2	g3	g0	g1	g2	g3	g0	g1	g2	g3	g0	g1	g2	g3	g0	g1	g2	g3
0	3.83	1.58	1.42	1.42	0	0	0	1	1	0	1	0	0	0	0	2					2	0	2	2
1	4.55	1.40	3.94	1.77	0	1	0	0					0	0	2	0					0	0	2	2
2	2.63	1.82	1.45	2.74	0	0	1	0	0	0	0	2	0	2	2	1	0	1	0	0	0	2	0	2
3	4.76	1.05	4.65	2.80	0	0	0	2					0	0	0	2	0	0	0	1	2	0	2	0
4	2.09	2.63	1.80	2.41	0	2	0	0	0	2	0	0	0	2	2	0	1	0	0	0				
5	2.18	4.86	1.22	1.93	1	0	0	0								0	2	0	0					
6	2.80	1.21	2.84	4.22	?	?	?	?					0	0	0	2	0	2	1	0	0	1	2	0
7	1.40	3.82	3.61	3.72	?	?	?	?					0	2	0	0					0	0	0	2
8	3.86	3.57	2.04	1.05	?	?	?	?	0	1	0	0			0	2	2				0	0	2	0
9	2.62	2.59	1.62	1.66	?	?	?	?					0	1	0	0	0	2	0	1	0	2	2	1
10	2.07	1.39	1.29	2.23	?	?	?	?	0	2	1	0	2	0	0	0	2	2	2	0	0	0	0	2

From Table 5.1 we see that student 0 ($sid = 0$) has already taken (and rated) 6 courses ($cid = 0, 1, 2, 3, 4, 5$) but has not taken the rest of the courses, so our task would be to recommend courses that this student could take in order either to increase his/her GAs in general (on average) or to increase a specific GA. These recommendations should take into account ratings of similar students to student 0 (collaborative filtering approach).

The CF algorithm utilizes the $C \times (G \times S)$ matrix (say R), like the one in Table 5.1. The first step is to declare a similarity measure between students in the matrix R . The similarity metric based on previous findings [10] is chosen to be Pearson Correlation and is provided by the Equation 5.2 for any two students a and b .

$$sim(a, b) = \frac{\sum_{g \in G_a \cap G_b} (R_{a,g} - \tilde{R}_a) \cdot (R_{b,g} - \tilde{R}_b)}{\sqrt{\sum_{g \in G_a \cap G_b} (R_{a,g} - \tilde{R}_a)^2} \cdot \sqrt{\sum_{g \in G_a \cap G_b} (R_{b,g} - \tilde{R}_b)^2}} \quad (5.2)$$

where:

- g represents an item which in our case is a pair of a Graduate Attribute and a course
- G_a represents the courses of student a (and similarly for b),

- $R_{a,g}$ represents the rating of student a for item g (and similarly for b)
- $G_a \cap G_b$ represents the common courses of students a and b ,
- \tilde{R}_a is the average rating of user a

The next step would be to find the neighbourhood of the active student which will define the set of students that will be used in order to generate predictions. Results have shown [48] that two techniques can effectively determine how many students will be included in the active student neighbourhood: Correlation thresholding and best n -neighbours with common courses threshold (direct application of the k -nearest neighbours algorithm). In the first case, we simply select all neighbours whose absolute correlation to the active student is higher than the value of the given threshold and include them in his/her neighbourhood. In the second case, we do not simply pick the best n correlates, but we ask that those students selected and the active student have rated a common number of course competences (in order to guarantee that a high correlation between two students is based on a decent number of courses) . In our experiments we applied Correlation Thresholding for a series of different correlation thresholds and we applied the second method (Nearest Neighbours) for different values of neighbourhood size n .

Finally, predictions for ratings is based on a weighted sum of ratings given to each item by similar students to our target student. This formula is given by the Equation 5.3.

$$P_{m,i} = \tilde{R}_m + \frac{\sum_{j \in N_u^K(m)} sim(m,j) \cdot (R_{j,i} - \tilde{R}_j)}{\sum_{j \in N_u^K(m)} |sim(m,j)|} \quad (5.3)$$

where:

- m is the active student,
- u is the neighbourhood of students close to m ,
- $N_u^K(m)$ is the K -most similar students to m

- $sim(m, j)$ is provided by Equation 5.2,
- $R_{j,i}$ is the rating of student j on item i ,
- \tilde{R}_m is the mean rating of student m ,
- \tilde{R}_j is the mean rating of student j .

One specific characteristic of this problem (recommending courses based on ratings) is that the most recent assessment has more value. For example assuming student S_1 is similar to student S_2 , the rating of student S_2 in the last semester is more valuable to make a recommendation based on, compared to a rating in a few semesters ago. To consider this feature in our algorithm we have given a higher weight to the more recent semester. The formula for this time based collaborative filtering algorithm is given in Equation 5.4.

$$P_{m,i} = \tilde{R}_m + \frac{\sum_{j \in N_u^K(m)} sim(m, j) \cdot e^{t_{R_{j,i}}} \cdot (R_{j,i} - \tilde{R}_j)}{\sum_{j \in N_u^K(m)} |sim(m, j) \cdot e^{t_{R_{j,i}}}|} \quad (5.4)$$

where:

- $t_{R_{j,i}}$ represents the semester in which rating $R_{j,i}$ is given.
- e is a constant (Euler's number)

At each semester, by having the previous ratings of students, we can predict the rating they would give to each pair (C,GA) of course and Graduate Attribute; i.e. how would the student assess the impact of a course on a Graduate Attribute. Then, we can recommend the courses with the maximum expected ratings. This recommendation can be targeted on a specific Graduate Attribute or the average of expected ratings for all the GAs.

5.4 Experiments and Evaluation

For the evaluation of the algorithm we use the synthetic data that we discussed before. At each semester we can predict the ratings of students on each of the




Figure 5.5: Mean squared error of rating predictor algorithms in each semester

possible pairs of course and GA. Then we would look at the exact value of the rating in the data and calculate the error of the prediction. For the baseline we use the CF algorithm provided in Equation 5.3 and we compare it with the other alternative provided in Equation 5.4.

Figure 5.5 shows the mean square error (MSE) of the prediction at each semester. This shows that in our targeted application considering the time factor can significantly improve the results.

In the results shown in Figure 5.5, the active user neighbourhood used for making predictions is the set of all students. By considering all the users in making the prediction we may achieve better results; however this is not an efficient method in terms of performance. As discussed in Section 5.3, there exist two techniques which can be used for determining the active user neighbourhood, Correlation thresholding and K best neighbours. In K best neighbours, for each targeted user, we rank all the other students based on their similarity, and we use the K top ones instead of all the users. In correlation thresholding, instead of having a fixed number of students, we use all the students which have a similarity to the target user higher than a specific threshold.

Figure 5.6 and Figure 5.7 show the MSE of the time based CF algorithm using different similarity thresholds along side the percentage of users used for making the prediction. As we increase the threshold, the error of the algorithm




Figure 5.6: Results of correlation thresholding for prediction at semester 7

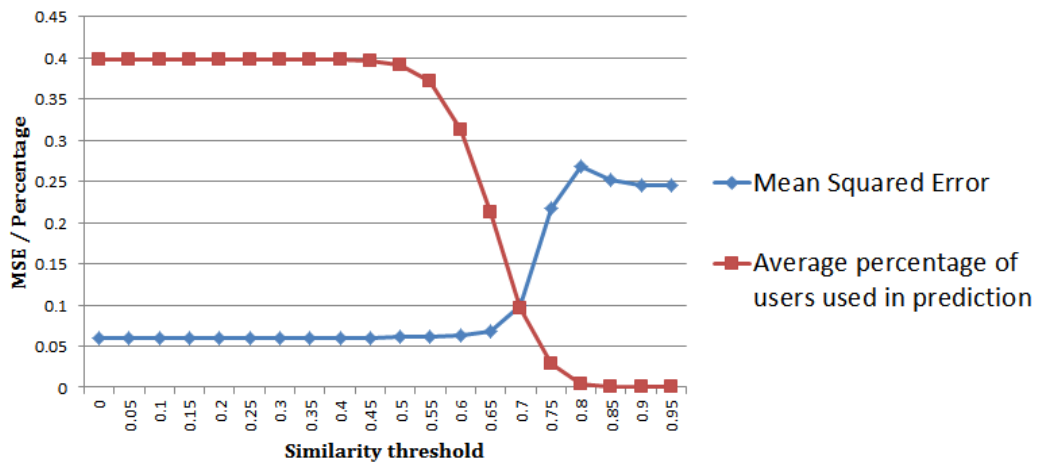


Figure 5.7: Results of correlation thresholding for prediction at semester 9

increases but at the same time less users are used for making the prediction which makes the algorithm faster. As a result we can choose to use a threshold such as 0.7 to decrease the number of users in the active user neighbourhood with a small increase in the error.

We have done the same experiment using K best neighbours methods. In this case instead of filtering the neighbourhood of the active user by a threshold we will choose the K most similar users. Figure 5.8 shows the effect of K on the MSE for the evaluation done in different semesters. The results show that the effective number of users in the neighbourhood can be between 10 to 15 students. This also matches with the result of similarity thresholding, as the threshold of 0.7 in our experiments corresponds to 10% of users, which is about 10 users. This shows that in terms of limiting the active neighbourhood, these two methods have similar results for our application.

5.5 Future Work

The main direction for the future of this recommender system can be the implementation based on the real data. The real data brings up some new challenges as well, which needs to be addressed in future works.

One of the challenges is the compatibility with the program plans of the institution. For example a recommender system should consider the prerequisites of courses in each program, since it is not appropriate to recommend courses which are not available to the student at the current time. Considering this issue, new options will be available for the logic of the algorithm. As an example, the algorithm may need to recommend a course which does not have a high value just to meet the prerequisites of a valuable course.

The other challenge with the implementation of the algorithm might be the scalability of the algorithm. To have a reasonable response time for making recommendations to a high number of students we may need to include new techniques. For example we can cluster similar students together using clustering algorithms. Then for making a recommendation we would only look into one cluster of similar students instead of all the students. Based on our




Figure 5.8: Results of K best neighbours method on prediction

results of correlation thresholding, we see that it is possible to look at a smaller set of students without sacrificing accuracy.

In future work we can also consider the performance of our algorithm, as collaborative filtering algorithms perform with less accuracy in the beginning when the data is still scarce. In the context of recommender systems this issue is known as the cold start problem. One possible option for improvement is the use of content based algorithms in the first few semesters and then over time giving more and more weight to the CF algorithm.

Chapter 6

Summary and Conclusions

In this project we designed and implemented an online assessment tool for assessment of Graduating Attributes (GAs) for University of Alberta. This tool has been implemented as a Moodle plugin, for compatibility with the course management system used in the University of Alberta. The purpose of this tool is the gathering of online assessments of students and instructors regarding the GAs. Then we studied the possible applications of assessing Graduating Attributes. We chose one of the possible applications (course recommender system) and implemented it as one of the useful applications. This recommender system will recommend courses with the purpose of improving expertise in Graduating Attributes. For our experiment we have generated some synthetic data based on the data model of GAAT, and we tested our algorithm with this data. The results shows that the course recommender system based on the Graduating Attributes is capable of recommending courses with the purpose of improving each or all of the GAs. Also, this application increases the engagement of students with the process of data gathering, as it promises the students better recommendations if they provide more accurate data.

Aside from the course recommender system, we suggested a few more applications in Chapter 4 which facilitates each of the stakeholders involved in our project, namely, instructors, administrators and employers.

The use of Graduating Attributes to its full potential has a long way to go as it needs the engagement of institution members including students, instructors

and administrators. But, doing so can bring some new values to educational institutions and perhaps even enhance the future of students after graduation.

Bibliography

- [1] Graduate attributes at the university of alberta. 2011. <http://www.provost.ualberta.ca/en/~media/provost/Documents/Information/GraduateAttributes.pdf>.
- [2] Rakesh Agrawal, Sreenivas Gollapudi, Anitha Kannan, and Krishnaram Kenthapadi. Study navigator: An algorithmically generated aid for learning from electronic textbooks. *JEDM-Journal of Educational Data Mining*, 6(1):53–75, 2014.
- [3] Marwah Alaofi and Grace Rumantir. Personalisation of generic library search results using student enrolment information. *JEDM-Journal of Educational Data Mining*, 7(3):68–88, 2015.
- [4] Heidi Andrade and Ying Du. Student responses to criteriareferenced self-assessment. *Assessment and evaluation in higher education*, 32(2):159–181, 2007.
- [5] Bahareh Azarnoush, Jennifer M. Bekki, George C. Runger, Bianca L. Bernstein, and Robert K. Atkinson. Toward a framework for learner segmentation. *JEDM-Journal of Educational Data Mining*, 5(2):102–126, 2013.
- [6] Ryan SJD Baker and Kalina Yacef. The state of educational data mining in 2009: A review and future visions. *JEDM-Journal of Educational Data Mining*, 1(1):3–17, 2009.
- [7] Trudy W. Banta and Catherine A. Palomba. *Assessment essentials: Planning, implementing, and improving assessment in higher education*. John Wiley and Sons, 2014.
- [8] J. Bowden, G. Hart, B. King, K. Trigwell, and O. Watts. Generic capabilities of atn university graduates. canberra: Australian government department of education, training and youth affairs. *Retrieved January*, 1:2011, 2000.
- [9] Javier Bravo and Alvaro Ortigosa. Detecting symptoms of low performance using production rules. *International Working Group on Educational Data Mining*, 2009.
- [10] John S Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 43–52. Morgan Kaufmann Publishers Inc., 1998.

- [11] J. Bruce Burke. Criteria for describing and assessing competency based programs. 1974.
- [12] Konstantina Chrysafiadi and Maria Virvou. Student modeling approaches: A literature review for the last decade. *Expert Systems with Applications*, 40(11):4715–4729, 2013.
- [13] Mihaela Cocea and Stephan Weibelzahl. Can log files analysis estimate learners’ level of motivation? 2006.
- [14] Gerben W. Dekker, Mykola Pechenizkiy, and Jan M. Vleeshouwers. Predicting students drop out: A case study. *International Working Group on Educational Data Mining*, 2009.
- [15] Naeimeh Delavari, Somnuk Phon-Amnuaisuk, and Mohammad Reza Beikzadeh. Data mining application in higher learning institutions. *Informatics in Education-International Journal*, 7:31–54, 2008.
- [16] Stanley Elam. Performance based teacher education. what is the state of the art?. 1971.
- [17] April Galyardt and Ilya Goldin. Move your lamp post: Recent data reflects learner knowledge better than older data. *JEDM-Journal of Educational Data Mining*, 7(2):83–108, 2015.
- [18] Enrique García, Cristóbal Romero, Sebastián Ventura, and Carlos De Castro. An architecture for making recommendations to courseware authors using association rule mining and collaborative filtering. *User Modeling and User-Adapted Interaction*, 19(1-2):99–132, 2009.
- [19] Norman E. Gronlund. *Assessment of student achievement*. ERIC, 1998.
- [20] Jiangang Hao, Zhan Shu, and Alina von Davier. Analyzing process data from game/scenario-based tasks: An edit distance approach. *JEDM-Journal of Educational Data Mining*, 7(1):33–50, 2015.
- [21] Jason M. Harley, Gregory J. Trevors, and Roger Azevedo. Clustering and profiling students according to their interactions with an intelligent tutoring system fostering self-regulated learning. *JEDM-Journal of Educational Data Mining*, 5(1):104–146, 2013.
- [22] Mohamed Osman Hegazi and Mazahir Abdelrhman Abugroon. The state of the art on educational data mining in higher education.
- [23] Tai-Chang Hsia, An-Jin Shie, and Li-Chen Chen. Course planning of extension education to meet market demand by using data mining techniques—an example of chinkuo technology university in taiwan. *Expert Systems with Applications*, 34(1):596–602, 2008.
- [24] Chiao-Tzu Huang, Wen-Tsann Lin, Shen-Tsu Wang, and Wen-Shan Wang. Planning of educational training courses by data mining: Using china motor corporation as an example. *Expert Systems with Applications*, 36(3):7199–7209, 2009.
- [25] Donald Ipperciel and Samira ElAtia. Assessing graduate attributes: Building a criteria-based competency model. *International Journal of Higher Education*, 3(3):p27, 2014.

- [26] John S. Kinnebrew, Kirk M. Loretz, and Gautam Biswas. A contextualized, differential sequence mining method to derive students' learning behavior patterns. *JEDM-Journal of Educational Data Mining*, 5(1):190–219, 2013.
- [27] Jared E. Knowles. Of needles and haystacks: Building an accurate statewide dropout early warning system in wisconsin. *JEDM-Journal of Educational Data Mining*, 7(3):18–67, 2015.
- [28] Chun-Hsiung Lee, Gwo-Guang Lee, and Yungho Leu. Application of automatically constructed concept map of learning to conceptual diagnosis of e-learning. *Expert Systems with Applications*, 36(2):1675–1684, 2009.
- [29] Ioanna Lykourantzou, Ioannis Giannoukos, Vassilis Nikolopoulos, George Mparadis, and Vassili Loumos. Dropout prediction in e-learning courses through the combination of machine learning techniques. *Computers & Education*, 53(3):950–965, 2009.
- [30] Collin Lynch, Kevin Ashley, Vincent Aleven, and Niels Pinkwart. Defining ill-defined domains; a literature survey. In *Proceedings of the workshop on intelligent tutoring systems for ill-defined domains at the 8th international conference on intelligent tutoring systems*, pages 1–10, 2006.
- [31] Leah P. Macfadyen and Shane Dawson. Mining lms data to develop an “early warning system” for educators: A proof of concept. *Computers & Education*, 54(2):588–599, 2010.
- [32] Aditi Mallavarapu, Leilah Lyons, Tia Shelley, and Brian Slattery. Developing computational methods to measure and track learners' spatial reasoning in an open-ended simulation. *JEDM-Journal of Educational Data Mining*, 7(2):49–82, 2015.
- [33] L. Dee Miller, Leen-Kiat Soh, Ashok Samal, Kevin Kupzyk, and Gwen Nugent. A comparison of educational statistics and data mining approaches to identify characteristics that impact online learning. *JEDM-Journal of Educational Data Mining*, 7(3):117–150, 2015.
- [34] Gali Naveh, Dorit Tubin, and Nava Pliskin. Student satisfaction with learning management systems: a lens of critical success factors. *Technology, Pedagogy and Education*, 21(3):337–350, 2012.
- [35] Joseph D. Novak and Alberto J. Cañas. The theory underlying concept maps and how to construct and use them. 2008. <http://cmap.ihmc.us/docs/theory-of-concept-maps>.
- [36] Michael P. O'Mahony and Barry Smyth. A recommender system for online course enrolment: an initial study. In *Proceedings of the 2007 ACM conference on Recommender systems*, pages 133–136. ACM, 2007.
- [37] Alejandro Peña-Ayala. *Educational Data Mining: Applications and Trends*, volume 524. Springer, 2013.
- [38] James W. Popham. *Transformative Assessment*. Association for Supervision and Curriculum Development. 1703 North Beauregard Street, Alexandria, VA 22311-1714., April 2008 2008. Source type: books; Object type: Book; CSAUnique: ED509137; AccNum: ED509137; DocISBN: 9781416606673.

- [39] R. Rallo, M. Gisbert, and J. Salinas. Using data mining and social networks to analyze the structure and content of educative on-line communities. *analysis*, 468(472):473, 1999.
- [40] Christophe Reffay and Thierry Chanier. *How social network analysis can help to measure cohesion in collaborative distance-learning*, pages 343–352. Designing for change in networked learning environments. Springer, 2003.
- [41] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. GroupLens: an open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pages 175–186. ACM, 1994.
- [42] Pablo Reyes and Pierre Tchounikine. Mining learning groups’ activities in forum-type tools. In *Proceedings of the 2005 conference on Computer support for collaborative learning: learning 2005: the next 10 years!*, pages 509–513. International Society of the Learning Sciences, 2005.
- [43] Cristobal Romero and Sebastian Ventura. Data mining in education. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 3(1):12–27, 2013.
- [44] Cristobal Romero, Sebastian Ventura, Mykola Pechenizkiy, and Ryan SJD Baker. *Handbook of educational data mining*. CRC Press, 2010.
- [45] Cristobal Romero, Amelia Zafra, Jose María Luna, and Sebastián Ventura. Association rule mining using genetic programming to provide feedback to instructors from multiplechoice quiz data. *Expert Systems*, 30(2):162–172, 2013.
- [46] Cristóbal Romero and Sebastián Ventura. Educational data mining: a review of the state of the art. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 40(6):601–618, 2010.
- [47] Jennifer L. Sabourin, Jonathan P. Rowe, Bradford W. Mott, and James C. Lester. Considering alternate futures to classify off-task behavior as emotion self-regulation: A supervised learning approach. *JEDM-Journal of Educational Data Mining*, 5(1):9–38, 2013.
- [48] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Analysis of recommendation algorithms for e-commerce. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 158–167. ACM, 2000.
- [49] John A. Self. Bypassing the intractable problem of student modelling.
- [50] Changjie Tang, Rynson WH Lau, Qing Li, Huabei Yin, Tong Li, and Danny Kilis. Personalized courseware construction based on web data mining. In *Web Information Systems Engineering, 2000. Proceedings of the First International Conference on*, volume 2, pages 204–211. IEEE, 2000.
- [51] U.S.DepartmentofEducation. Competency-based learning or personalized learning. <http://www.ed.gov/oii-news/competency-based-learning-or-personalized-learning>.

- [52] Cesar Vialardi, Javier Bravo Agapito, Leila Shila Shafti, and Alvaro Ortigosa. *Recommendation in higher education using data mining techniques*. Barnes, T., Desmarais, M., Romero, C., Ventura, S., 2009.
- [53] Victor CX Wang. *Handbook of research on education and technology in a changing society*. IGI Global, 2014.
- [54] Andrew Waters, Christoph Studer, and Richard Baraniuk. Collaboration-type identification in educational datasets. *JEDM-Journal of Educational Data Mining*, 6(1):28–52, 2014.
- [55] Judith Zimmermann, Kay H. Brodersen, Hans R. Heinemann, and Joachim M. Buhmann. A model-based approach to predicting graduate-level performance using indicators of undergraduate-level performance. *JEDM-Journal of Educational Data Mining*, 7(3):151–176, 2015.

Appendix A

List of information targeted by the Online Questionnaire

- How different are the students from different programs and departments in terms of GAs?
- How the year of study can affect students' self-evaluation in terms of GAs?
- Is there any correlation between the extracurricular activities and any of the GAs? And if there is which activities has correlation with which of GAs?
- Is there any correlation between the time students spend on studying and any of the GAs? And if there is how is this correlation and with which of GAs?
- Does living on campus or off campus affect the GAs? And if it does, how does it affect GAs?
- How different instructors from different disciplines May plan their courses in terms of GAs?