

Grade Prediction with Course and Student Specific Models

Agoritsa Polyzou^(✉) and George Karypis

University of Minnesota, Minneapolis, USA
polyz001@umn.edu

Abstract. The accurate estimation of students' grades in future courses is important as it can inform the selection of next term's courses and create personalized degree pathways to facilitate successful and timely graduation. This paper presents future-course grade predictions methods based on sparse linear models and low-rank matrix factorizations that are specific to each course or student-course tuple. These methods identify the predictive subsets of prior courses on a course-by-course basis and better address problems associated with the *not-missing-at-random* nature of the student-course historical grade data. The methods were evaluated on a dataset obtained from the University of Minnesota. This evaluation showed that the course specific models outperformed various competing schemes with the best performing scheme achieving a RMSE across the different courses of 0.632 vs 0.661 for the best competing method.

1 Introduction

Data mining and machine learning approaches are being increasingly used to analyze educational- and learning-related datasets towards understanding how students learn and improving learning outcomes. This has led to the development of various approaches for modeling and predicting the success or failure of students in completing specific tasks in the context of intelligent tutoring systems [9, 12, 15, 16, 18, 19], building intelligent “early warning systems” that monitor the students' performance during the term [1, 3], predicting how well the students will perform by analyzing their activities with the learning management system (e.g., Moodle) [8, 11, 17], and predicting students' term and final GPA [2, 13, 14].

Our work focuses on developing methods that utilize historical student-course grade information to accurately estimate how well students will perform (as measured by their grade) on courses that they have not yet taken. Being able to accurately estimate students' grades in future courses is important as it can be used by them (and/or their academic advisers) to identify the appropriate set of courses to take during the next term, and create personalized degree pathways that enable them to successfully and effectively acquire the required knowledge to complete their studies in a timely fashion.

Existing approaches for predicting a student’s grade in a future course [4, 6, 7] rely on neighborhood-based collaborative filtering methods [10]. Despite their relative simplicity, the estimations obtained by these methods are reasonably accurate indicating that there is sufficient information in the historical student-course grade data to make the estimation problem feasible.

In this paper we improve upon these methods by developing various future-course grade prediction methods that utilize approaches based on sparse linear models and low-rank matrix factorizations. These methods rely entirely on the performance that the students achieved in previously taken courses. A unique aspect of many of our methods is that their associated models are either specific to each course or specific to each student-course tuple. This allows them to identify and utilize the relevant information from the prior courses that are associated with the grade for each course and better address problems associated with the *not-missing-at-random* nature of the student-course historical grade data. We experimentally evaluated the performance of our methods on a dataset obtained from the University of Minnesota that contained historical grades that span 12.5 years. Our results showed that the course specific models outperformed various competing schemes and that the best performing scheme, which is based on course-specific regression, achieves a RMSE across the different courses of 0.632 whereas the best competing method achieves an RMSE of 0.661.

The remainder of the paper is organized as follows. Section 2 introduces the notation and definitions used. Section 3 describes the methods developed and Sect. 4 provides information about the experimental design. Section 5 presents an extensive experimental evaluation of the methods and compares them against existing approaches. Finally, Sect. 6 provides some concluding remarks.

2 Definitions and Notations

Throughout the paper, bold lowercase letters will denote column vectors (e.g., \mathbf{y}) and bold uppercase letters will denote matrices (e.g., \mathbf{G}). Individual elements will be denoted using subscripts (e.g., for a vector y_i , and for a matrix $g_{s,c}$). A single subscript on a matrix will denote its corresponding row. The sets will be represented by calligraphic letters.

The historical student-course grade information will be represented by a sparse matrix $\mathbf{G} \in \mathbb{R}^{n \times m}$, where n and m are the number of students and courses, respectively, and $g_{i,j}$ is the grade in the range of $[0,4]$ that student i achieved in course j . If a student has not taken a course, the corresponding entry will be missing. The course and student whose grades need to be predicted will be called *target course* and *target student*, respectively.

3 Methods

In this section we describe various classes of methods that we developed for predicting the grade that a student will obtain on a course that he/she has not yet taken.

3.1 Course-Specific Regression (CSR)

Undergraduate degree programs are structured in such a way that courses taken by students provide the necessary knowledge and skills for them to do well in future courses. As a result, the performance that a student achieved in a subset of the earlier courses can be used to predict how well he/she will perform in future courses. Motivated by this, we developed a grade prediction method, called *course-specific regression* (CSR) that predicts the grade that a student will achieve in a specific course as a sparse linear combination of the grades that the student obtained in past courses.

In order to estimate the CSR model for course c , we extract from the overall student-course matrix \mathbf{G} the set of rows corresponding to the students that have taken c . For each of these students (rows), we keep only the grades that correspond to courses taken prior to course c . Let $\mathbf{G}^c \in \mathbb{R}^{n_c \times m}$ be the matrix representing that extracted information, where n_c is the number of students that took course c . In addition, let $\mathbf{y}^c \in \mathbb{R}^{n_c}$ be the grades that the students in \mathbf{G}^c obtained in course c (the y_i^c is the grade corresponding the student in the i th row of \mathbf{G}^c). Given this, the CSR model $\mathbf{w}^c \in \mathbb{R}_+^m$ for c is estimated as:

$$\underset{\mathbf{w}^c \geq 0}{\text{minimize}} \quad \|\mathbf{y}^c - \mathbf{1}w_0^c - \mathbf{G}^c\mathbf{w}^c\|_2^2 + \lambda_1 \|\mathbf{w}^c\|_2^2 + \lambda_2 \|\mathbf{w}^c\|_1, \quad (1)$$

where w_0^c is a bias term, $\mathbf{1} \in \mathbb{R}^{n_c}$ is a vector of ones and λ_1, λ_2 are regularization parameters to control overfitting and promote sparsity. The model is non-negative because we assume that prior courses can only provide knowledge to future courses. The individual weights of \mathbf{w}^c indicate how much each prior course contributes in the prediction and represent a measure of the importance of the prior course within the context of the estimated model. Using this model, the grade that a student will obtain in course c is estimated as

$$\hat{y}^c = w_0^c + \mathbf{s}^T \mathbf{w}^c, \quad (2)$$

where $\mathbf{s} \in \mathbb{R}^m$ is the vector of the student's grades in the courses he/she has taken so far.

We found that by centering each student's grades around his/hers GPA leads to more accurate predictions (see Sect. 5.1). In this approach, prior to estimating the model using Eq. 1, we first subtract from each $g_{i,j}^c$ grade the GPA of each student (GPA is calculated based on the information in \mathbf{G}^c). This centers the data for each student and takes into consideration a notion of student bias as it predicts the performance with respect to the current state of a student. Note that in the case of GPA-centered data, we remove the non-negativity constraint on \mathbf{w}^c . We will refer to this model as the CSR-RC (Row Centered) model.

3.2 Student-Specific Regression (SSR)

Depending on the major, the structure of different undergraduate degree programs can be different. Some degree programs have limited flexibility as to the

set of courses that a student has to take and at which point in their studies they can take them (i.e., specific semester). Other degree programs are considerably more flexible and are structured around a fairly small number of core courses and a large number of elective courses.

For the latter type of degree programs, a drawback of the CSR method is that it requires the same linear regression model to be applied to all students. However, given that the set of prior courses taken by students in such flexible degree programs can be quite different, a single linear model can fail to capture the various prior course combinations. In fact, there can be cases in which many of the most important courses that were identified by the CSR model were simply not taken by some students, even though these students have acquired the necessary knowledge and skills by taking a different set of courses. To address this limitation, we developed a different method, called *student-specific regression* (SSR), which estimates course-specific linear regression models that are also specific to each student.

The student specific model is derived by creating a student-course specific grade matrix $\mathbf{G}^{s,c}$ for each target student s and each target course c from the \mathbf{G}^c matrix used in CSR method. $\mathbf{G}^{s,c}$ is created in two steps. First, we eliminate from \mathbf{G}^c any grades for courses that were not taken by the target student. Second, we eliminate from \mathbf{G}^c the rows that correspond to students that have not taken a sufficient number of courses that are in common with the target student s . Specifically, if \mathcal{C}_s and \mathcal{C}_i are the set of courses for student s and i respectively, we compute the overlap ratio (OR) = $|\mathcal{C}_s \cap \mathcal{C}_i|/|\mathcal{C}_s|$ and if $\text{OR} < t$, then student i is not included in $\mathbf{G}^{s,c}$. The value of t is a parameter of the SSR method and high values ensure that the set of students forming $\mathbf{G}^{s,c}$ have taken many courses in common with s and have followed similar degree plans. Given $\mathbf{G}^{s,c}$, the SSR method proceeds to estimate the model using Eq. 1 (with $\mathbf{G}^{s,c}$ replacing \mathbf{G}^c), and uses Eq. 2 for prediction.

3.3 Methods Based on Matrix Factorization

Low rank matrix factorization (MF) approaches have been shown to be very effective for accurately estimating ratings in the context of recommender systems [10]. These approaches can be directly applied to the problem of predicting the grade that a student will achieve on a particular course by treating the student-course grade matrix \mathbf{G} as the user-item rating matrix.

The use of such MF-based approaches for grade prediction is postulated on the fact that there is a low dimensional latent feature space that can jointly represent both students and courses. Given the nature of the domain, this latent space can correspond to the space of knowledge components. Each course vector is the set of components associated with a course and each student vector represents the student's level of knowledge across these knowledge components.

By applying the common approaches of MF-based rating prediction to the problem of grade prediction, the grade that student i will obtain on course j is estimated to be

$$\hat{g}_{i,j} = \mu + sb_i + cb_j + \mathbf{p}_i \mathbf{q}_j^T, \quad (3)$$

where μ is a global bias term, sb_i and cb_j are the student and course bias terms, respectively, and p_i and q_j are the latent representations for student i and course j , respectively. The parameters of the MF method ($\mu, \mathbf{s}\mathbf{b} \in \mathbb{R}^n, \mathbf{c}\mathbf{b} \in \mathbb{R}^m, \mathbf{P} \in \mathbb{R}^{n \times l}$, and $\mathbf{Q} \in \mathbb{R}^{m \times l}$) are estimated following a matrix completion approach that considers only the observed entries in \mathbf{G} as

$$\begin{aligned} \underset{\mu, \mathbf{s}\mathbf{b}, \mathbf{c}\mathbf{b}, \mathbf{P}, \mathbf{Q}}{\text{minimize}} \quad & \sum_{g_{i,j} \in \mathbf{G}} (g_{i,j} - \mu - sb_i - cb_j - \mathbf{p}_i \mathbf{q}_j^T)^2 + \lambda (\|\mathbf{P}\|_F^2 + \|\mathbf{Q}\|_F^2 \\ & + \|\mathbf{s}\mathbf{b}\|_2^2 + \|\mathbf{c}\mathbf{b}\|_2^2), \end{aligned} \quad (4)$$

where λ is a regularization parameter and l is the dimensionality of the latent space, which is a parameter to this method.

The accurate recovery of the low rank model (when such a model exists) from a set of partial observations depends on having a sufficient number of observed entries, and on these entries be randomly sampled from the entries of the target matrix \mathbf{G} [5]. However, in the context of student grade data, the set of courses that students take is not a random subset of the courses being offered as they need to satisfy their degree program requirements. As a result, such an MF approach may lead to suboptimal prediction performance.

In order to address this problem we developed a *course specific matrix factorization* (CSMF) approach that estimates an MF model for each course by utilizing a course specific subset of the data that is denser (in terms of the number of observed entries and the dimensions of the matrix). As a result, it contains a larger number of random by sampled subsets of sufficient size.

Given a course c and a set of students \mathcal{S}^c for which we need to estimate their grade for c (i.e., the students in \mathcal{S}^c have not taken this course yet), the data that CSMF utilizes are the following: (i) the students and grades of the \mathbf{G}^c matrix and \mathbf{y}^c vector of the CSR method (Sect. 3.1), (ii) the students in \mathcal{S}^c and their grades. This data is used to form a matrix $\mathbf{X}^c \in \mathbb{R}^{(n_c + n_t) \times (m_c + 1)}$, where n_c is the number of students in \mathbf{G}^c , $n_t = |\mathcal{S}^c|$, and m_c is the number of distinct courses that have at least one grade in \mathbf{G}^c or \mathcal{S}^c . The values stored in \mathbf{X}^c are the grades that exist in \mathbf{G}^c and \mathcal{S}^c . The last column of \mathbf{X}^c stores the grades \mathbf{y}^c for the course c that were obtained from the students in \mathbf{G}^c . Thus, \mathbf{X}^c contains all the prior grades associated with the students who have already taken course c and the students for which we need to have their grade on c predicted. Matrix \mathbf{X}^c is then used in place of matrix \mathbf{G} in Eq. 4 to estimate the parameters of the CSMF method, which are then used to predict the missing entries of the last column of \mathbf{X}^c , which are the grades that need to be predicted.

4 Experimental Design

4.1 Dataset

The student-course-grade dataset that we used in our experiments was obtained from the University of Minnesota which has a very flexible degree program.

Table 1. Statistics for course-specific datasets.

Prior courses	5	7	9
Average number of students in training set	270	232	212
Average number of students in test set	22	21	20
Average number of prior courses	141	141	145
Average number of grades	3,872	3,663	3,663
Courses predicted	92	90	80
Grades predicted	2,088	1,959	1,666

It contains the students that have been part of the Computer Science and Engineering (CSE) and Electrical and Computer Engineering (ECE) programs from Fall of 2002 to Spring of 2014. Both of these degree programs are part of the College of Science & Engineering (CS&E) in which students have to take a common set of core science courses during the first 2–3 semesters. We removed from the dataset any courses that are not part of those offered by CS&E departments, as these correspond to various liberal arts and physical education courses, which are taken by few students and in general do not count towards degree requirements. Furthermore, we eliminated any courses that were taken as pass/fail. The initial grades were in the A–F scale, which was converted to the 4–0 scale using the standard letter-grade to GPA conversion. The resulting dataset consists of 2,949 students, 2,556 different courses, and 76,748 student-course grades.

We used this dataset to assess the performance of the different methods for the task of predicting the grades that the students will obtain in the last semester (i.e., the most recent semester for which we have data). For this reason, the dataset was further split into two parts, one containing the students that are still *active*, i.e., have taken courses in the last semester (D_{active}) and one that contains the remaining students ($D_{inactive}$). D_{active} contains 876 students, 19,089 grades, out of which 3,427 grades are for the 475 distinct classes taken in the last semester. $D_{inactive}$ contains 2,073 students and 57,659 grades.

These datasets were used to derive various training and testing datasets for the different methods that we developed. Specifically, for the CSR method we extracted the course specific training and testing datasets as follows. For each course c that was offered in the last semester, we extracted course-specific training and testing sets ($D_{train}^{c, \geq k}$ and $D_{test}^{c, \geq k}$) by selecting from $D_{inactive}$ and D_{active} , respectively, the students that have taken c , and prior to taken c , they also took at least k other courses. The reason that these datasets were parametrized with respect to k is because we wanted to assess how the methods perform when different amount of historical student performance information is available. In our experiments we used k in the set $\{5, 7, 9\}$. That information will create the grade matrix \mathbf{G}^c , where $g_{i,j}^c$ is the grade of the i th student on the j th course from the training set $D_{train}^{c, \geq k}$. Table 1 shows various statistics about the various course-specific datasets for different values of k .




Fig. 1. Statistics of the datasets used in SSR w.r.t overlap ratio.

For the CSMF method, the training dataset for course c was obtained by combining $D_{train}^{c, \geq k}$ and $D_{test}^{c, \geq k}$ into a single matrix after removing the grades that the target students achieved in course c .

For the MF method, the matrix is constructed as the union of the sets $D_{train}^{c, \geq k}$ and $D_{test}^{c, \geq k}$ for every course to be predicted after removing the grades that the active students achieved in the courses we want to predict. We formulated the dataset in this way in order to provide the same information for training and testing to all our models.

In the SSR, the grade matrix $\mathbf{G}^{s,c}$ is created by selecting from $D_{train}^{c, \geq k}$ the set of courses that were also taken by student s and the set of students whose OR with s is at least t . Figure 1 shows some statistics about these datasets as a function of t .

Finally, we did not consider the models that have less than 20 students in their corresponding dataset, as we consider them to have too few training instances for reliable estimation.

4.2 Competing Methods

In our experiments, we compared our methods with the following competing approaches.

1. **BiasOnly.** We only took into consideration local and global bias to predict the students' grades. These biases were estimated using Eq. 4 when $l = 0$.
2. **Student-Based Collaborative Filtering (SBCF).** This method implements the approach described in [4]. For a target course c , every student i is represented by a vector formed with his/hers grades in courses taken prior to c . The vector of a target student s is compared against the vectors of the other students that have taken course c with the Pearson's correlation coefficient. We select the students with positive similarity to perform grade prediction

for s in c according to:

$$\hat{g}_{s,c} = \bar{g}_s + \frac{\min(r, nbr)}{r} \frac{\sum_{i=1}^{nbr} (g_{i,c} - \bar{g}_i) sim_{s,i}}{\sum_{i=1}^{nbr} sim_{s,i}}, \quad (5)$$

where nbr is the number of students selected, r is a confidence lower limit for significance weighting, \bar{g}_i is the average grade of the student prior taking c , and $sim_{s,i}$ represents the similarity of target student s with i .

4.3 Parameters and Model Selection

For CSR, we let λ_1 take values from 0 to 40 in increments of 2.5 and λ_2 from 0 to 50 in increments of 2.5. For SSR, we let λ_1 take values from 0 to 10 in increments of 1 and λ_2 from 0 to 14 in increments of 2. For MF and CSMF, we let λ take values from 0 to 6 in increments of 0.05. For SSR, the range of the tested values for overlap ratio is 0.3 to 1, in increments of 0.04. For MF and CSMF methods we tested the number of latent dimensions with the values 2, 5 and 8.

As we could not use cross validation for the SSR, we did not apply it for any regression model, in order to be fair with our comparisons. The best models are selected based on their performance on the test set. For MF based approaches, we used the semester before the target semester to estimate and select the best parameters.

4.4 Evaluation Methodology and Performance Metrics

We evaluated the performance of the different approaches by using them to predict the grades for the last semester in our dataset using the data from the previous semester for training.

We assessed the performance using the root mean square error (RMSE) between the actual grades and the predicted ones. Since the courses whose grades are predicted have different number of students, we computed two RMSE-based metrics. The first is the overall RMSE in which all the grades across the different courses were pooled together, and the second is the average RMSE obtained by averaging the RMSE values for each course. We will denote the first by RMSE and the second as AvgRMSE.

5 Experimental Results


5.1 Course-Specific Regression

Table 2 shows the performance achieved by the CSR and CSR-RC models when trained using the three different datasets discussed in Sect. 4.1. These results show that among the two models, CSR-RC, which operates on the GPA-centered grades leads to considerably lower errors both in terms of RMSE and AvgRMSE.

Table 2. The performance achieved by Linear Course-Specific Regression.

	RMSE			AvgRMSE		
	5	7	9	5	7	9
CSR	0.751	0.761	0.779	0.757	0.785	0.762
CSR-RC	0.634	0.632	0.632	0.585	0.579	0.543

The performance of the models trained on the different datasets were evaluated on the $D_{test}^{\geq 9}$ test set, which is the common subset among their respective test sets.

**Fig. 2.** The performance achieved by the SSR model w.r.t. overlap ratio.

In terms of the sensitivity of their performance on the amount of historical information that was available when estimating these models (i.e., the minimum number of prior courses), we can see that for CSR-RC, the RMSE performance of the models does not change significantly; though the AvgRMSE performance improves when going from five to nine prior courses. This indicates that training sets with more number of prior courses tend to help smaller courses.

5.2 Student-Specific Regression

As one of the parameters for this problem was the overlap ratio between the courses of the target student and other students, Fig. 2 presents the behavior of the model's RMSE (left) and AvgRMSE (right) as we vary the overlap ratio for $D_{test}^{c, \geq 5}(k = 5)$, $D_{test}^{c, \geq 7}(k = 7)$ and $D_{test}^{c, \geq 9}(k = 9)$. When the overlap ratio is increased, the selected students have more courses in common with the target user and that results to better performance. In order to compare the performance of SSR against CSR-RC, Fig. 3 shows the RMSE of the best CSR-RC and SSR models. The RMSE values were computed as the subsets of the test set that was predicted by both models. If the overlap ratio is more than 0.8, then SSR is more accurate. However, the capability of this method to predict courses is very low, i.e., we can predict 50 % less courses than the CSR model for $k = 9$ when the overlap ratio is more than 0.8, because there are not as many students that had followed the same degree plan as the selected student.




Fig. 3. RMSE of SSR model compared to the CSR-RC w.r.t. overlap ratio for the case of 9 prior courses. The performance for other choices of prior courses is very similar.

5.3 Methods Based on Matrix Factorization

The performance of the methods based on matrix factorization is shown in Table 3 for various number of latent factors. Besides the MF and CSMF schemes that were described in Sect. 3.3, this table also shows results for a method labeled “MF-GB”, which is derived from the MF scheme by eliminating the global bias term (μ) of Eq. 4. These results show that CSMF leads to lower RMSE values when there are more than nine prior courses per student, which confirms that by building matrix factorization models on smaller but denser course-specific sub-matrices, we can derive low-rank models that lead to more accurate matrix completion. Even for the case with more than five prior courses, if we focus on denser models, the majority of courses are predicted better by CSMF* than by the best model, MF-GB. In terms of the number of latent factors, we can see that in most cases, the best performance is achieved with small number of latent factors. This should not be surprising, as the average number of grades per student is low, which does not support a large number of latent factors.

5.4 Comparison with other methods

Table 4 compares the performance of the baseline approaches described in Sect. 4.2 (BiasOnly and SBCF) with the best-performing course-specific regression method (CSR-RC), and the best CSMF method (two latent factors). In addition, the results labeled “CSMF*” correspond to those obtained by CSMF in which the best-performing number of latent factors for each course can be different and was selected based on their performance on the validation set (10% of the training data). CSR-RC and CSMF lead to RMSE and AvgRMSE values that are substantially better than either BiasOnly or SBCF. In terms of the methods that we developed, we see that CSR-RC consistently outperforms CSMF, suggesting that sparse linear regression methods are better than those based on matrix factorization for this setting. Finally, comparing the performance of CSMF* against CSMF, we see that even though the former achieved

Table 3. The performance achieved by the methods based on matrix factorization model w.r.t. the number of prior courses and the number of latent factors.

Prior courses	Latent Factors		MF	MF-GB	CSMF	CSMF*
5	2	RMSE	0.662	0.661	0.683	0.676
	5		0.666	0.667	0.682	0.682
	8		0.667	0.672	0.679	0.676
	2	AvgRMSE	0.597	0.581	0.648	0.645
	5		0.603	0.569	0.643	0.647
	8		0.604	0.596	0.645	0.644
7	2	RMSE	0.667	0.671	0.684	0.679
	5		0.673	0.675	0.680	0.677
	8		0.676	0.681	0.681	0.676
	2	AvgRMSE	0.590	0.598	0.641	0.643
	5		0.603	0.607	0.638	0.640
	8		0.604	0.610	0.637	0.640
9	2	RMSE	0.675	0.684	0.683	0.671
	5		0.677	0.687	0.676	0.672
	8		0.681	0.692	0.677	0.674
	2	AvgRMSE	0.581	0.600	0.653	0.648
	5		0.582	0.607	0.645	0.646
	8		0.579	0.599	0.648	0.647

Table 4. Comparison of the performance achieved from our methods with the competing approaches.

Prior courses	RMSE			AvgRMSE		
	5	7	9	5	7	9
BiasOnly			0.728			0.687
SBCF			0.677			0.675
CSR-RC	0.634	0.632	0.632	0.585	0.579	0.543
CSMF	0.679	0.680	0.676	0.645	0.638	0.645
CSMF*	0.676	0.676	0.671	0.644	0.640	0.648

The performance of the models trained on the different datasets were evaluated on the $D_{test}^{\geq 9}$ test set, which is the common subset among their respective test sets.

better performance, the difference is not very large, which suggests that CSMF's performance is more consistent across its different model parameters.

6 Conclusions

In this paper, we presented two course-specific approaches based on linear regression and matrix factorization that perform better than existing approaches based on traditional methods. This suggests that focusing on a course specific subset of the data can result in more accurate predictions. A student-course specific approach was also developed but its accuracy in grade prediction is limited by the diverse nature of degree plans. The course-specific regression was the one with the best results compared to any other method tested.

Aknowledgements. This work was supported in part by NSF (IIS-0905220, OCI-1048018, CNS-1162405, IIS-1247632, IIP-1414153, IIS-1447788) and the Digital Technology Center at the University of Minnesota. Access to research and computing facilities was provided by the Digital Technology Center and the Minnesota Supercomputing Institute. <http://www.msi.umn.edu>.

References

1. Starfish: Earlyalert. <http://www.starfishsolutions.com/home/student-success-solutions/>. Accessed 4 October 2015
2. Al-Barrak, M.A., Al-Razgan, M.: Predicting students final GPA using decision trees: a case study. *Int. J. Inf. Educ. Technol.* **6**(7), 528 (2016)
3. Arnold, K.E., Pistilli, M.D.: Course signals at purdue: using learning analytics to increase student success. In: *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge*, pp. 267–270. ACM (2012)
4. Bydžovská, H.: Are collaborative filtering methods suitable for student performance prediction? In: Pereira, F., Machado, P., Costa, E., Cardoso, A. (eds.) *EPIA 2015. LNCS*, vol. 9273, pp. 425–430. Springer, Heidelberg (2015)
5. Chen, Y., Bhojanapalli, S., Sanghavi, S., Ward, R.: Coherent matrix completion. *arXiv preprint* (2013). [arXiv:1306.2979](https://arxiv.org/abs/1306.2979)
6. Denley, T.: Course recommendation system and method. <http://www.google.com/patents/US20130011821>. Accessed 4 October 2015
7. Denley, T.: Austin peay state university: degree compass. *EDUCAUSE Review Online* (2012). <http://www.educause.edu/ero/article/austin-peay-state-university-degree-compass>
8. Elbadrawy, A., Studham, R.S., Karypis, G.: Collaborative multi-regression models for predicting students' performance in course activities. In: *Proceedings of the Fifth International Conference on Learning Analytics And Knowledge*, pp. 103–107. ACM (2015)
9. Hwang, C.S., Su, Y.C.: Unified clustering locality preserving matrix factorization for student performance prediction. *IAENG Int. J. Comput. Sci.* **42**(3), 245–253 (2015)
10. Kantor, P.B., Rokach, L., Ricci, F., Shapira, B.: *Recommender Systems Handbook*. Springer, New York (2011)

11. Luo, J., Sorour, E., Goda, K., Mine, T.: Predicting student grade based on free-style comments using word2vec and ANN by considering prediction results obtained in consecutive lessons, pp. 396–399, June 2015
12. McKay, T., Miller, K., Tritz, J.: What to do with actionable intelligence: E 2 coach as an intervention engine. In: Proceedings of the 2nd International Conference on Learning Analytics and Knowledge, pp. 88–91. ACM (2012)
13. Ogunde, A., Ajibade, D.: A data mining system for predicting university students' graduation grades using ID3 decision tree algorithm. *J. Comput. Sci. Inf. Technol.* **2**(1), 21–46 (2014)
14. Osmanbegović, E., Suljić, M.: Data mining approach for predicting student performance. *Econ. Rev.* **10**(1), 3–12 (2012)
15. Pardos, Z.A., Heffernan, N.T.: Using HMMs and bagged decision trees to leverage rich features of user and skill from an intelligent tutoring system dataset. *J. Mach. Learn. Res. W & CP* (2010)
16. Romero, C., Ventura, S., Espejo, P.G., Hervás, C.: Data mining algorithms to classify students. In: Educational Data Mining 2008 (2008)
17. Sorour, S.E., Mine, T., Goda, K., Hirokawa, S.: A predictive model to evaluate student performance. *J. Inf. Process.* **23**(2), 192–201 (2015)
18. Thai-Nghe, N., Drumond, L., Horváth, T., Schmidt-Thieme, L.: Using factorization machines for student modeling. In: UMAP Workshops (2012)
19. Toscher, A., Jahrer, M.: Collaborative filtering applied to educational data mining. In: KDD Cup (2010)