

Impacts of Students' Academic-Performance Trajectories on their Final Academic Success

Mr. Shahab Boumi, University of Central Florida

Shahab Boumi is a Ph.D. candidate in the department of Industrial Engineering and Management Systems (IEMS) at the University of Central Florida. His main research focus on educational data mining and machine learning.

Prof. Adan Ernesto Vela

Impacts of Students Academic-Performance Trajectories on Final Academic Success

Abstract

Many studies in the field of education analytics have identified student grade point averages (GPA) as an important indicator and predictor of students' final academic outcomes (graduate or halt). And while semester-to-semester fluctuations in GPA are considered normal, significant changes in academic performance may warrant more thorough investigation and consideration, particularly with regards to final academic outcomes. However, such an approach is challenging due to the difficulties of representing complex academic trajectories over an academic career. In this study, we apply a Hidden Markov Model (HMM) to provide a standard and intuitive classification over students' academic-performance levels, which leads to a compact representation of academic-performance trajectories. Next, we explore the relationship between different academic-performance trajectories and their correspondence to final academic success. Based on student transcript data from University of Central Florida, our proposed HMM is trained using sequences of students' course grades for each semester. Through the HMM, our analysis follows the expected finding that higher academic performance levels correlate with lower halt rates. However, in this paper, we identify that there exist many scenarios in which both improving or worsening academic-performance trajectories actually correlate to higher graduation rates. This counter-intuitive finding is made possible through the proposed and developed HMM model.

Keywords: Hidden Markov model, student GPA analysis, academic outcomes, clustering

1 Introduction

Based on the National Student Clearinghouse Research Center, the six-year graduation rate for students who initiated their college education in 2012 is only 58%, with 42% of students either halting enrollment or taking longer than six years to graduate [1]. Halting college is broadly understood to impose irreversible mental, financial, and time losses to students [2].

To date, numerous reasons have been identified to explain why students choose to halt their education; they include: financial problems, lack of interest in studies, lagged behind in study progression, and inadequate information and guidance [3]. In many cases, the reasoning is related to academic performance; however, the determination to halt one's education is not a straight line between perceived poor grades and *halting*. As D. Shippee notes, poor grades may lead to the feeling of depression [4], or lead students to shift career goals [5]. Instead, low academic performance can be viewed as an underlying factor, predictor, or leading indicator of halting.

Along these lines, many studies have recognized GPA as a fair predictor or leading indicator of students persisting or halting their academic career [6, 7, 8]. Zahedi et al. [6] investigated the relationship between graduation rate and numerous variables, including gender, race, transfer status, major, number of terms registered, and cumulative GPA. Their findings suggest that cumulative GPA is one the most powerful measures for predicting graduation. In a similar study, Pappas et al. [7] showed that students' cumulative GPA is a significant factor in identifying students at the risk of dropout among computer science students. Although cumulative GPA is demonstrated to be a valuable indicator and predictor of a students' academic outcomes, the measure is a summary statistic. As such, cumulative GPA does not convey information of how a student's GPA evolved over time. Using cumulative GPA as a predictive measure essentially ignores underlying patterns that could convey meaningful status changes.

To better understand, model, and characterize students at-risk of halting, we believe it is necessary to mathematically model academic-performance trajectories. The desire to model academic performance trajectories is driven by the following research question:

- Does a decrease in academic performance proportionally increase the risk of halting?
- If a student improves their academic performance does their graduation rate match those students that always performed well?

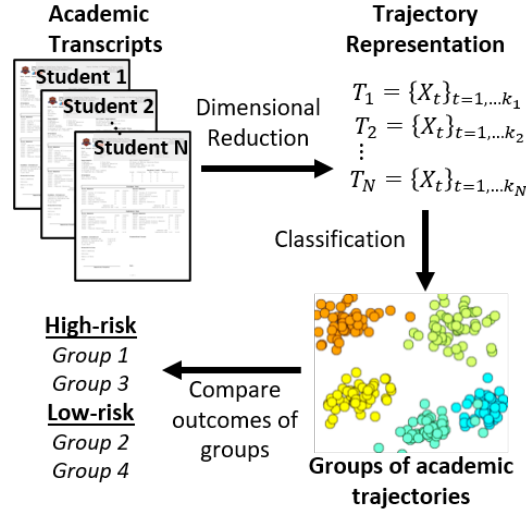
This study is aimed at analyzing students' academic-performance trajectories to see if these trajectories are meaningful in answering the research questions above. To address the research questions, we start by designing a Hidden Markov Model (HMM) to convert complex student academic records to a compact discrete representation we denote as an *academic-performance level*. The HMM is tuned using anonymized student records of first-time-in-college (FTIC) students at the University of Central Florida. After applying the HMM, the resulting compact representation of the academic-performance trajectories is grouped into meaningful clusters for analysis. Groups are compared and contrasted to understand how academic-performance trajectories correlate to halting. Accordingly, the primary contributions of this paper are: **(1)** a technique for generating a compact representation of a student's academic-performance trajectory; and **(2)** an improved understanding of the relationship between academic-performance trajectories and final academic outcomes.

The remainder of the paper is organized as follows: in Section 2 we provide a more detailed problem statement, with Section 3 detailing our proposed approach. Next, we provide our findings in Section 4. Finally, the paper is concluded in Section 5.

2 Problem Statement

We consider the problem of mathematically modeling the trajectory dynamics of a student's academic performance. The goal of addressing this problem is to identify a compact trajectory representation of academic performance that enables the classification and clustering of undergraduate students into meaningful groups. These groups will then be compared and contrasted to explore the relationship between academic performance trajectories and final educational outcomes, specifically halting. A pictorial representation of the overall problem and methodology is provided in Figure 1.

Figure 1: Representation of overall problem and solution methodology.



Here, we start with the notional definition of an academic-performance trajectory to be a temporal sequence of states indicating a student’s academic performance. A frequently encountered category of academic performance trajectories are student transcripts; each semester corresponds to a time-step t in a trajectory \hat{T}_i , while courses are taken, and grades received each semester correspond to trajectory states, \hat{X}_t . The challenge with working such data is the dimensionality of the representation - specifically all the combinations of courses and grades in \hat{X}_t .

The goal of this paper is to identify a mapping $\hat{T}_i \rightarrow T_i$ that translates an academic performance trajectory into a lower-dimensional representation T_i , where the state X_t at each semester corresponds to an academic-performance level. As an example, consider the case in which each student transcript is first converted to a sequence of grade counts given by the number of A’s, B’s, C’s, D’s/F’s (combined), and withdrawals they received in a semester. One such tuple, $(1,0,1,1,2)$, would indicate that in the corresponding semester, a student received the following marks: one A, no B’s, one C, one D or F, and two withdrawals. Over an academic career, a semester-to-semester sequence might look like:

$$(1, 0, 1, 1, 2), (2, 1, 1, 1, 0), (0, 3, 1, 0, 0), \dots, (3, 1, 0, 0, 0).$$

A mapping of the full trajectory to a compact trajectory representation might be $L_3, L_2, L_2, \dots, L_1$, where each L_i in the time-ordered sequence corresponds to a specified academic-performance level (e.g. L_1 is low, L_2 is middle, and L_3 is high).

Unlike prior studies, we explicitly assume the grades a student earns exist within a random process whereby students’ course grades are not deterministic. As an analogy, in an imaginary multiverse, a student could take the same course one-hundred times and earn different scores each time. Seemingly random processes or disturbances that might affect a student’s scores, and ultimately final grade, can be due to external factors (e.g., illness the week of an exam) or internal factors (e.g., grading policies) related to the course. Accordingly, instead of each academic-performance level, L_i being a straight-forward mapping from the semester GPA, random processes and

disturbances that create a noisy GPA signal must be filtered to estimate the academic-performance level.

3 Methodology

In order to map full academic trajectories to a more compact representation, we develop and apply a Hidden Markov Model (HMM). Application of HMMs to model students' academic behaviors and performance has been documented in several studies [9, 10, 11, 12, 13]. For example, in work by Boumi and Vela [14] the authors use an HMM to investigate the relationship between students' academic performance and enrollment strategies (i.e., the sequence of full-time and part-time enrollment statuses). Through the application of an HMM, students are classified and clustering into three categories based on their enrollment strategy – full-time, part-time, and mixed enrollment strategy. Once clustered, the categorizations are demonstrated to correlate with various measures of academic performance (e.g., graduation rate, halting, GPA). In another study by Kaser et al. [15], the authors apply an HMM to exam the impact of students' exploration strategies on learning. Their findings suggest that exploration strategy is important to learning outcomes. Using a similar framework to these studies, we develop and apply an HMM to model and understand the impact of academic-performance trajectories and final educational outcomes.


In this section, we provide an overview of the Hidden Markov Model used in analyzing students' *academic-performance trajectories*. Followed is a discussion of the problem formulation and prerequisite steps to encoded academic records and develop the HMM structure. Lastly, we provide the relevant parameters of the trained HMM model.

3.1 Representing Academic Performance using a Hidden Markov Model

As depicted in Figure 2, an HMM represents the dynamics of a system as it moves between operating states or modes (e.g., Modes 1, 2, and 3 in the figure). When operating within a state or mode, the system generates a mode-related output O_i at each time-step. For the problem under consideration here, the modes correspond to the academic-performance level of a student, and the observations refer to the student grade counts in any given semester. Because of the randomness assumption stated in Section 2, the grade counts observed might not map directly to the expected academic-performance level of a student, and as such, we say the mode or academic-performance level of a student is not directly observable from their grades. Instead, through a series of grade observations, we can estimate or infer the most likely academic-performance level of the student in any semester given the likelihood of such a sequence of mode transitions (i.e., trajectory).

Each Hidden Markov Model is initialized by a set of modes $Q = \{q_1, q_2, \dots, q_N\}$. Unlike any other types of Markov models in which the modes are observable, in HMMs, modes are hidden; that is, it is unknown at which mode the system dwells at any particular moment. After developing an HMM, one application of an HMM is then to estimate a variable-length sequence of the hidden modes estimates given a series of observations, $O = o_1, o_2, \dots, o_T$, that are generated from the modes. Within this framework, each HMM is defined by the following parameters: (1) $A = [a_{i,j}] \in \mathcal{R}^{N \times N}$, a transition matrix containing the single-step probabilities by which a system moves between modes; (2) $B = b_i(o_t)$, an emission matrix describing the probability of

Figure 2: Representation of a simple Hidden Markov Model



generating each observation o_t when the system is in mode q_i ; and (3) π_0 , the initial probability distribution for the system at its start.

The first step to build an HMM, known as the **LEARNING STEP**, is to estimate the optimal values for the parameters $\lambda = (A, B, \pi_0)$. We use the Baum-Welch algorithm, an expectation-maximization algorithm, to estimate the optimal parameter-set λ^* . This algorithm takes initial guesses for each of the parameters and iteratively improves the estimations of the parameters by computing the likelihood of any sequence of observation given λ . After obtaining the optimal λ^* , the next step is to predict the sequence of hidden modes for a given observation sequence. This step, known as **DECODING STEP** uses the Viterbi algorithm, which takes observation sequences and λ^* as input and returns estimated hidden modes as the output.

In our case, the first challenge with utilizing an HMM is the dimensionality of the observable outputs, i.e., grade counts. As described in Section 2, each semester, we represent a student's academic performance according to a tuple with five elements, in which the first, second, and third elements are integers indicating the number of courses with A, B, and C grades. The fourth and fifth elements are variables indicating the number of D or F grades, and how many courses the student has withdrawn. Under the assertion that the grade counts earned each semester are dependent (e.g., the number of A's earned is correlated to the number of B's, C's, etc.), it is reasonable to consider each unique grade-tuple to be a possible observation. Within the data-set under consideration, over 687 unique grade combinations in a semester have been observed. The most common grade combination was (4,0,0,0,0), which occurred 5.8% of the time. Meanwhile, there are over 20 grade combinations that occur only once, e.g. (1,1,1,0,4).

Given the large number of grade combinations, the wide distribution spread, and the relative sparseness in which many cases occur, it is necessary to perform a partial reduction in the dimensionality of the output space. Instead of each output being a tuple of unconstrained counts, elements of the grade-tuple are capped. The grade-tuple is adjusted to contain the number of A's, B's, and C's up to a maximum value of 3. Furthermore, instead of storing the number of D/F and W grades, these values are converted to binary variables. In this new representation the original grade-tuple (5, 0, 2, 0, 3) for a student that had registered for 10 classes would be converted to (3, 0, 2, 0, 1). Now, because students do not typically register for more than 5 classes, and more

Table 1: Cumulative GPA for the four academic-performance levels obtained by the hierarchical clustering method

Level number	Cumulative GPA
1	$CGPA \leq 2.5$
2	$2.5 < CGPA \leq 3$
3	$3.0 < CGPA \leq 3.5$
4	$3.5 < CGPA \leq 4.0$

so, because there is a strong correlation between grades, the impact of clipping is believed to be limited. After applying the new representation, the total number of unique observations is reduced to 216 grade-tuples.

Within the HMM framework, the next step is to specify a predetermined number of modes for the model; again, each mode corresponds to an academic-performance level. The inclusion of each additional state increases the required number of parameters used to define the model (i.e., size of λ^*); it is critical to balance over-fitting (too many parameters) and model coarseness (too few parameters). For a model with N and M outputs, the total number of parameters within an HMM to be tuned is $N(N + M - 2) + N$; the tuned parameters correspond to a probability transition matrix, the emission probability for each grade tuple when operating in a state, and the initial state distribution. Through a combination of managing the number of modes and the number of grade-tuples represented, we can ensure there is sufficient data to tune the parameters.

Initially, we classify students based on their last semester's cumulative GPAs (CGPA) into eight academic-performance levels. The first level corresponds to students with a CGPA less than 2.0; level 2 includes students with a CGPA between 2.0 and 2.25, level 3 for students with CGPA between 2.25 and 2.5, and so on for other levels. Next, for each academic-performance level, the distribution over grades including A, B, C, DF, and W was computed. Since the grade distributions observed over adjacent academic-performance levels were quite close, a round of aggregation over academic-performance levels was initiated using their pairwise distance. Based on [16], the distance between two Multinomial distribution, Δ , is computed by the following equation:

$$4 \sin^2(\Delta/2) = (\sqrt{p_A} - \sqrt{p_A}) + \dots + (\sqrt{p_W} - \sqrt{p_W})$$

In the above equation, p_A through p_W corresponds to the probability of grade A through grade W in a given academic-performance level. Given the pairwise distances between academic-performance levels and using the hierarchical clustering method, the eight different academic-performance levels were further grouped into four main levels. We consider the resulting four academic-performance levels as the hidden states for the proposed HMM. The cumulative GPAs associated with these states are shown in Table 1.

Based on an HMM model with 4 academic-performance levels and 216 possible grade-tuples, it is possible to provide examples of the desired input and output of the HMM model; see Table 2. In the table, student number 1 has enrolled in four semesters and has four grade-tuples recorded in his academic history. In the first semester, he has no courses with A, B, or C grades as the first three digits are all zero. The fourth and fifth digits in the grade-tuple indicate that the student has grades DF and W in his first semester. His grade-tuple for the rest of the semesters is explained in

Table 2: Example students’ grade-tuple sequence and corresponding estimated academic-performance level sequence

Student	Grade-tuple Seq.	Academic-per. Seq.
1	00011,10100,00211,00011	1,1,1,1
2	11110,11200,21001,13000	2,2,2,2
3	22000,32000,30100	3,3,3
4	30000,31000,30000	4,4,4
5	30000,30000,00010,00010	4,4,1,1

the same manner. The estimated academic-performance level for this student in all enrolled semesters is 1. As we see in the table, while students number 1 to 4 has a consistent academic-performance level during their academic careers, the academic-performance level for student number 5 has changed from 4 to 1 in his third semester.

3.2 Initialization and Training

For this particular problem, the process for initializing and training the HMM is non-trivial, especially given the total number of parameters that define the emission distributions (i.e., 4 states x 216 possible grade-tuple outputs).

First, we describe the process for providing an initial guess for the smaller 4x4 transition matrix, which has 12 linearly independent parameters (out of 16 total). To identify an initial guess for HMM transition matrix, we analyzed students’ transcripts to estimate how their GPA moves between the four levels identified in Table 1 – this is equivalent to assuming a students GPA directly maps to an academic performance level (removing the randomness assertion stated earlier). The matrix below provides the resulting initial transition matrix (percentage scale):

$$\text{Initial guess for } A = \begin{bmatrix} 81.81 & 17.10 & 1.02 & 0.07 \\ 7.90 & 78.43 & 13.52 & 0.15 \\ 0.74 & 8.54 & 84.44 & 6.28 \\ 0.34 & 1.00 & 10.32 & 88.34 \end{bmatrix}$$

As an initial guess for the vector π , we consider the distribution over students’ academic-performance in their first semester. The matrix below shows the initial guess for π_0 (percentage scale):

$$\text{Initial guess for } \pi_0 = [20.14 \quad 19.07 \quad 33.50 \quad 27.29]$$

As we explained earlier, an HMM observation in this study for a student in a given semester is defined as a tuple with five elements. Elements 1 to 3 store the number of courses with grades A, B, and C, respectively, and elements 4 and 5 are binary variables, determining if the student has D/F and W grades in that semester. In order to find an initial guess for the emission matrix, we supposed that elements of the observed grade-tuples are independent of each other, each having a Poisson distribution (in the tuning process, the independence assumption between grade counts is removed). As such, for each academic-performance level, we can compute the probability of

observing each grade-tuple, i.e., the grades counts in each tuple, using a joint but separable Poisson probability density function. Since there is a total of 216 different grade-tuple combinations in our data set, the initial emission matrix has four rows (corresponding to the four hidden states) and 216 columns (corresponding to the 216 observations). Due to space limitations, instead of showing the initial emission matrix with 864 parameters, we have computed the expected values (EV) of the counts for each grade across the 216 possible observations. The results are shown in the following matrix:

$$\text{Initial EV(grade)} = \begin{bmatrix} 0.331 & 0.676 & 0.718 & 0.304 & 0.154 \\ 0.702 & 1.058 & 0.726 & 0.179 & 0.100 \\ 1.130 & 1.077 & 0.401 & 0.070 & 0.060 \\ 1.359 & 0.572 & 0.075 & 0.009 & 0.025 \end{bmatrix}$$

In the above matrix, rows 1 through 4 correspond to academic-performance levels 1 to 4, while columns 1 to 5 provide the expected value for the number of courses with an A, B, C, D/F, and W grade each semester. For example, for students in academic-performance level 3 (the first row), in each semester, in the expectation, students earn 1.13 A's, 1.077 B's, and .401 C's.

After initializing HMM parameters, the next step is to estimate their optimal value using the Baum-Welch algorithm. The optimal estimates, obtained after 20 iterations for algorithm convergence, are listed below:

$$A = \begin{bmatrix} 86.11 & 10.93 & 2.44 & 0.52 \\ 8.40 & 80.25 & 11.26 & 0.09 \\ 0.29 & 7.58 & 84.98 & 7.15 \\ 0.02 & 0.02 & 5.64 & 94.32 \end{bmatrix}$$

$$\text{EV(grade)} = \begin{bmatrix} 0.302 & 0.658 & 0.862 & 0.608 & 0.291 \\ 0.789 & 1.470 & 0.902 & 0.210 & 0.098 \\ 1.777 & 1.416 & 0.297 & 0.036 & 0.054 \\ 2.561 & 0.451 & 0.023 & 0.002 & 0.020 \end{bmatrix}$$

$$\pi_0 = [12.82 \quad 37.35 \quad 35.79 \quad 14.04]$$

$$\pi = [11.70 \quad 18.19 \quad 30.44 \quad 39.67]$$

. Again, instead of providing the probability of observing each of the 864 possible grade-tuples, the expectation of earned grades is reported.

4 Analysis

The trained values of the HMM reported in Section 3 are of value in understanding academic-performance trajectories. Based on the estimated transition matrix A , most students maintain their academic-performance level (mode) from one semester to the next. Looking at the diagonal of the matrix, 86.11% of students in academic-performance level 1 in semester t will remain at the same academic-performance level in semester $t + 1$. These probabilities for students in academic-performance levels 2, 3, and 4 are 80.25%, 84.98%, and 94.32%. Based on these values, we are able to assert, with exception to performance-level 1, that increasing

Table 3: Distribution over students' academic-performance level at UCF

Academic-per. level	Percentage
Level 1	11.75%
Level 2	14.42%
Level 3	19.10%
Level 4	11.43%
Other	43.30%

performance-levels are increasingly stable, so that a student in performance-level 4 is more likely to remain in the same performance-level as compared to a student in performance-level 2 or 3.

Matrix $EV(\text{grade})$ provides the expected values for grade counts using the estimated emission matrix. As illustrated in the matrix, students in academic-performance level 4 on average have more courses with A grades compared to the students in academic-performance level 1 ($2.561 > 0.302$). Technically, while W grades are not used to compute semester GPAs, the last column of the matrix $EV(\text{grade})$ indicates there is a significant relationship between W grades and students' academic-performance level. Based on the columns, students with academic-performance level 1 are more likely to withdraw from courses in any given semester compared to students in other academic-performance levels ($0.291 > 0.098 > 0.054 > 0.02$).

Matrix π_0 provides an estimate of the distribution of academic-performance levels for students' first semester. Based on the matrix, most students start their academic careers at academic-performance levels 2 and 3 ($37.35\%+35.79\%$). The matrix π , corresponding to the stationary probability distribution of the transition matrix A, reports the average probability students are in each academic-performance level, regardless of semester. The Comparison between π_0 and π provides an indication of how the distribution of students' academic-performance evolves over time. For instance, it illustrates that students are more likely to be in academic-performance level 4 in later semesters compared to the first semester ($39.67\% > 14.04\%$). Furthermore, the probability of having academic-performance level 1 in the first semester is more than the same probability for the later semesters ($12.82\% > 11.70\%$). There are two potential reasons for the change in the distribution of academic-performance levels over time: first, since there is a significant correlation between students' first GPA and retention [17], students with low first semester GPA may halt, which causes the probability of students being in academic-performance level 1 to decreased for later semesters. Second, improvement in students' GPA over time could be another reason for the mentioned change in the academic-performance level distribution. Table 3 depicts the distribution over students' academic-performance level at UCF. Levels 1 to 4 correspond to the students who maintain a consistent academic-performance level over all enrolled semesters. The group *Other*, corresponds to the students who have a transition between academic-performance levels during their academic career at UCF. The synthetic student number 5 in Table 2 is one such example for this category of students. This student has an academic-performance level of 4 in his first two semesters and then switch to academic-performance level 1 in his third semester. As reported in the figure, 43.3% of students change academic-performance levels during their academic careers.

Table 4: Distribution over switching type for students with one switch in academic-performance level

Switch	Ratio	Switch	Ratio
level 1 →level 2	2.4%	level 3 →level 1	0.2%
level 1 →level 3	1.0%	level 3 →level 2	12.9%
level 1 →level 4	0.5%	level 3 →level 4	27.7%
level 2 →level 1	22.3%	level 4 →level 1	0.0%
level 2 →level 3	27.8%	level 4 →level 2	0.0%
level 2 →level 4	0.0%	level 4 →level 3	5.2%

Table 5: Comparing halt rate between students with no switch in academic-performance level and one switch academic-performance level

Staying in	Halt rate	Switching	Halt rate
Level 1	97.4%	level 1 →level 2	9.2%
		level 1 →level 3	6.5%
		level 1 →level 4	5.1%
Level 2	43.7%	level 2 →level 1	73.6%
		level 2 →level 3	4.2%
		level 2 →level 4	–
Level 3	23.8%	level 3 →level 1	77.3%
		level 3 →level 2	8.7%
		level 3 →level 4	2.0%
Level 4	12.8%	level 4 →level 1	100%
		level 4 →level 2	–
		level 4 →level 3	3.2%

Of particular interest to this paper is the *Other* group, which consists of students whose academic-performance changes levels during their academic career. Among these students, 75.2% switch their academic-performance levels only once, 20.4% switch their academic-performance levels twice, and 4.5% change their academic-performance levels more than 2 times during their academic career at UCF. The distribution over switching type for students with one switch is summarized in Table 4. Based on the table, the more common switches between academic-performance levels corresponds to level 2 →level 3 (27.8%), level 3 →level 4 (27.7%), and level 2 →level 1 (22.3%). For students with one switch in their academic-performance level, 59.4% improve their academic-performance levels, while 40.6% worsen. For the group with two switches, the percentages of students ultimately improving and worsening their academic-performance level are 41.1% and 15.7%, respectively. Moreover, 43.2% of these students go back to their initial academic-performance level after two switches (for example, level 3 →level 2 →level 3).

Table 5 compares the halt rate for a subset of academic-performance level trajectories. Column *Staying in* corresponds to students who keep a consistent academic-performance level. As the corresponding halt rate column suggests, the higher academic performance level, the lower the halt rate (97.4% > 43.7% > 23.8% > 12.8%). Furthermore, we see that students that switch

academic-performance levels from level 1 to other levels have a substantially improved halt rate (from 97.4% to 9.2%, 6.5%, and 5.1% for levels 2, 3, and 4). Also, switching from any level to level 1 increases the halt rate significantly. For example, all students whose academic-performance levels have changed from 4 to 1 left school without a degree. This percentage for students from level 3 and 2 to level 1 are 73.6% and 77.3%. However, there are surprising results for students starting at academic-performance level 3, where we see that students who worsen to academic-performance level 2 have a lower halt rate than if than students that had remained at level 3 (23.8% > 8.7%). A similar pattern is observed for students whose academic-performance levels are changed from 4 to 3. When combined, this evidence implies that switching from a high academic-performance level to a lower academic-performance level does not necessarily result in a higher halt rate. Another surprising result emerges from the grouping and analysis of the academic trajectories that involve improving academic performance. For students that improve their academic performance, their halt rate is substantially lower than those students that maintained a consistent academic performance level. For example, the halt rate for students who transitioned from level 2 to level 3 is lower than those students who were always at level 3 (4.2% < 23.8%). Even students that transitioned from level 1 to level 2 have a lower halt rate than those students consistently in level 4.

5 Conclusion and Future Research

In this paper, the authors analyzed students' academic-performance trajectories at University of Central Florida and examined their final academic outcomes in terms of halt rates. Unlike traditional statistical methodologies, our proposed approach is able to provide a standard point of reference for comparing student's GPA both across their enrolled semesters and across all students. Our proposed HMM model takes the sequence of course grades over multiple semesters and returns the sequence of estimated academic-performance levels. The estimated HMM parameters illustrate that while W grades are not involved in computing students' GPA, there is a significant relationship between students' academic-performance level and withdrawing from courses. Also, by tracking students' transitions between four main academic-performance levels, we observed that a significant portion of students (43.30%), acquire a combination of academic-performance levels during their education. Furthermore, in analyzing and comparing the halt rate for students with consistent academic performance levels, evidence suggests that students who constantly maintain a low academic-performance level are more likely to leave school without a degree. Meanwhile, for students with a change in their academic-performance level, it was shown that switching from any academic-performance levels to level 1 increases the halt rate substantially. Most surprising, however, was that for some academic-performance trajectories, when switching from a high academic performance level to a lower level, the halt rate did not increase but rather decreased. In fact, the reduction in halt rate is even lower than if the student had remained at the same academic level. An investigation into the underlying reasons for such observations is left as a good avenue for future qualitative research.

As discussed earlier in this paper, one of the main challenges of HMMs correspond to the dimensionality of the observations. In order to tackle this problem, two different approaches were used in this chapter. First, an upper limit was considered on the elements of the grade-tuple which restricted number of A's, B's, and C's to the maximum of 3 and transformed the number of D/F

and W grades to binary variables. Second, since the probability of some grade combinations is zero, each unique grade-tuple was considered as one possible HMM observation. This assumption decreased the total number of unique observations to 216 (from 687). Although leveraging these problem-specific features helped reduce size of the problem that was addressed in this paper, in general such conditions may not always exist. The alternative approach to tackle this problem can be considered as a research question for future studies.

References

- [1] Doug Shapiro, Afet Dundar, Faye Huie, Phoebe Khasiala Wakhungu, Ayesha Bhimdiwala, and Sean Eric Wilson. Completing college: A national view of student completion rates—fall 2012 cohort (signature report no. 16). *National Student Clearinghouse*, 2018.
- [2] Pietro Garibaldi, Francesco Giavazzi, Andrea Ichino, and Enrico Rettore. College cost and time to complete a degree: Evidence from tuition discontinuities. *Review of Economics and Statistics*, 94(3):699–711, 2012.
- [3] Elisabeth Hovdhaugen and Per Olaf Aamodt. Learning environment: Relevant or not to students’ decision to leave university? *Quality in higher education*, 15(2):177–189, 2009.
- [4] Nathan D Shippee and Timothy J Owens. Gpa, depression, and drinking: A longitudinal comparison of high school boys and girls. *Sociological Perspectives*, 54(3):351–376, 2011.
- [5] Elva Hull-Blanks, Sharon E Robinson Kurpius, Christie Befort, Sonja Sollenberger, Megan Foley Nicpon, and Laura Huser. Career goals and retention-related factors among college freshmen. *Journal of Career Development*, 32(1):16–30, 2005.
- [6] Leila Zahedi, Stephanie J. Lunn, Samira Pouyanfar, Monique S. Ross, and Matthew W. Ohland. Leveraging machine-learning techniques to analyze computing persistence in undergraduate programs. In *2020 ASEE Virtual Annual Conference Content Access*, number 10.18260/1-2–34921, Virtual On line, June 2020. ASEE Conferences. <https://peer.asee.org/34921>.
- [7] Ilias O Pappas, Michail N Giannakos, and Letizia Jaccheri. Investigating factors influencing students’ intention to dropout computer science studies. In *Proceedings of the 2016 ACM Conference on Innovation and Technology in Computer Science Education*, pages 198–203, 2016.
- [8] Farshid Marbouti, Heidi A Diefes-Dux, and Krishna Madhavan. Models for early prediction of at-risk students in a course using standards-based grading. *Computers & Education*, 103: 1–15, 2016.
- [9] Manie Tadayon and Gregory J Pottie. Predicting student performance in an educational game using a hidden markov model. *IEEE Transactions on Education*, 2020.
- [10] Masun Homsni, Rania Lutfi, Rosa M Carro, and Barakat Ghias. A hidden markov model approach to predict students’ actions in an adaptive and intelligent web-based educational system. In *2008 3rd International Conference on Information and Communication Technologies: From Theory to Applications*, pages 1–6. IEEE, 2008.

- [11] Kristy Elizabeth Boyer, Robert Phillips, Amy Ingram, Eun Young Ha, Michael Wallis, Mladen Vouk, and James Lester. Investigating the relationship between dialogue structure and tutoring effectiveness: a hidden markov modeling approach. *International Journal of Artificial Intelligence in Education*, 21(1-2):65–81, 2011.
- [12] Mohammad H Falakmasir, José P González-Brenes, Geoffrey J Gordon, and Kristen E DiCerbo. A data-driven approach for inferring student proficiency from game activity logs. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*, pages 341–349. ACM, 2016.
- [13] Loc Nguyen. A new approach for modeling and discovering learning styles by using hidden markov model. *Global Journal of Human Social Science Linguistics & Education*, 13(4), 2013.
- [14] Shahab Boumi and Adan Vela. Application of hidden markov models to quantify the impact of enrollment patterns on student performance. *International Educational Data Mining Society*, 2019.
- [15] Tanja Käser, Nicole R Hallinen, and Daniel L Schwartz. Modeling exploration strategies to predict student performance within a learning environment and beyond. In *Proceedings of the seventh international learning analytics & knowledge conference*, pages 31–40, 2017.
- [16] Anil Bhattacharyya. On a measure of divergence between two multinomial populations. *Sankhyā: the indian journal of statistics*, pages 401–406, 1946.
- [17] Susan Gershenfeld, Denice Ward Hood, and Min Zhan. The role of first-semester gpa in predicting graduation rates of underrepresented students. *Journal of College Student Retention: Research, Theory & Practice*, 17(4):469–488, 2016.