

Variations in patterns of persistence

Adan Vela,¹ Jacquelyn J. Chini,² Alexander Baekey,² and Joseph Walsh²

¹*Department of Industrial Engineering and Management Systems,
University of Central Florida, 12800 Pegasus Drive, Orlando, Florida, USA*

²*Department of Physics, University of Central Florida, 4111 Libra Drive, Orlando, Florida, 32816*

Despite improved college accessibility, female students continue to be underrepresented in physics, mathematics, and computer science. More so, the gender gap widens from enrollment to graduation as female students transition out of the majors more frequently than men. While there exist several theoretical frameworks modeling persistence and retention, a common theme amongst them involves the process by which students respond to challenges. Accordingly, this research explores the responses of students to academic performance within their major and overall curriculum. In particular, this study makes use of mathematical models to explore differences in persistence at a large state institution according to academic performance within their major, overall technical curriculum, and non-technical course work. This research represents a contribution over prior works that limited consider to overall GPA when analyzing retention.

I. INTRODUCTION

Women continue to be underrepresented among computer science, physics, and mathematics bachelor's degree earners in the United States. In fact, within a ten-year span from 2004 to 2014 the percentage of bachelor's degrees awarded to women in the United States decreased in the fields of mathematics, statistics and physics, with a significant drop in computer science (CS) from 25.1% to 18.1% [1]. Low representation of female students within CS is also present at the University of Central Florida (UCF), a large public research university located in Orlando, Florida. Between 2008 and 2017 female students only accounted for 14% of all students who at one time declared themselves to be CS majors. In spite of increases in the proportion of female enrollments in CS during the same time period, female students are leaving the major at higher rates than their male counterparts resulting in low representation among CS bachelor's degree earners (see Table I). Of particular interest, however, is that first-time in college female students entering CS, leaving CS in good academic standing (above 2.0 GPA), and graduating in CS are on average equaling or outperforming male students (see Table II) – which appears contradictory to their relative attrition rate when compared to male students. Superficially, this mismatched performance-to-retention rate comparison between genders is consistent with another study regarding the relative performance of female students that leave engineering [2]. So while there exists known gender differences in early exposure to computer science [3], such discrepancies might not be present at UCF or might not significantly affect the early academic performance of females in the field.

The decision for students to persist within technical fields has been well studied. Research on both majority and underrepresented groups have indicated that academic performance measures, particularly grade point averages (GPA), serve as key indicators for predicting if students will continue within a science, technology, engineering, or math (STEM) major [4, 5]. Limited research however has attributed the decision process to leave a major according to absolute and rel-

ative academic performance within major courses (i.e. CS courses), courses within the greater technical curriculum (e.g. math, physics, chemistry), and courses outside the technical curriculum (e.g. general education courses). Key exceptions include [6–8], of which [8] hypothesizes differing grade sensitivities as an explanation for the reported gender imbalances within an undergraduate economics degree program. Given that female students associated with UCF's CS program equal or outperform male students in overall GPA, additional examination of the decision process to continue or exit the CS degree program is worthy of attention. Accordingly, the major contribution of this paper is a study of the academic performance factors influencing persistence of male and female students within CS at UCF within the first three semesters of enrollment.

TABLE I. Persistence within CS at UCF.

Status	Male Number (% by gender, % within gender)	Female Number (% by gender, % within gender)	Total Number (%)
Declared	4349 (86%, 100%)	704 (14%, 100%)	5053 (100%)
Enrolled	918 (85%, 21%)	164 (15%, 23%)	1082 (21%)
Grad.	774 (90%, 18%)	83 (10%, 12%)	857 (16%)
Attrited	2657 (85%, 61%)	457 (15%, 65%)	3114 (61%)

TABLE II. Average GPA of CS students at UCF.

Gender	High School	Exiting	Graduating
Male	3.79	2.93	3.26
Female	4.02	3.03	3.42

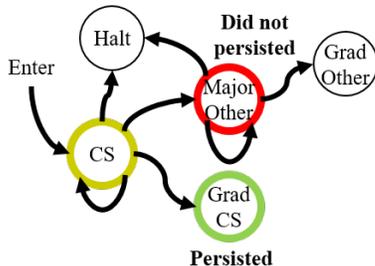


FIG. 1. State transition diagram tracking CS students through their academic career.

II. DATA AND METHODOLOGY

We have chosen to perform a comparative cluster analysis on male and female students in CS using academic student data. The intent is to extract academic performance characteristics that drive the observed persistence differences between male and female students. While focusing on CS, the future goal is understanding the phenomenon of high attrition rates for female students in the other majors of mathematics and physics. The cluster analysis considers students in good standing who have either graduated or attrited from CS. A visual representation of persistence is illustrated in Figure 1. As depicted in the figure, a student is declared to *not persist in CS* if having once enrolled in CS they are later enrolled in another major. A student is considered to persist in CS if they graduate with a CS degree. All other indeterminate cases are excluded from analysis, this includes students who have halted their education at the university and students still enrolled in CS that have yet to graduate.

The analysis is further restricted to students in good standing. Good standing is defined as maintaining a cumulative GPA over 2.0 in all courses, and a cumulative GPA above 2.0 in courses within the CS curriculum (this includes CS courses as well as math, physics, and other technical courses). The GPA restrictions mostly remove from consideration students placed on academic probation and students who are academically disqualified from the major or from the university. In these cases, the decision to leave CS is removed from the student and delegated to the university.

Student profiles within in the dataset are described according to gender and objective measures of academic performance, along with ultimate persistence outcomes. The academic performance measures include: high school GPA (hs-GPA), cumulative university GPA (GPA), cumulative GPA within CS courses (mGPA), cumulative GPA within the full technical curriculum (fcGPA), cumulative GPA in non-CS courses within the technical curriculum (cGPA), cumulative GPA in other courses (oGPA), number of withdrawals within major (nmW) and within curriculum (ncW), and the number of semesters enrolled within CS thus far. Other traits and flags describing race/ethnicity, first-generation college student, financial aid, and family income are also available, however,

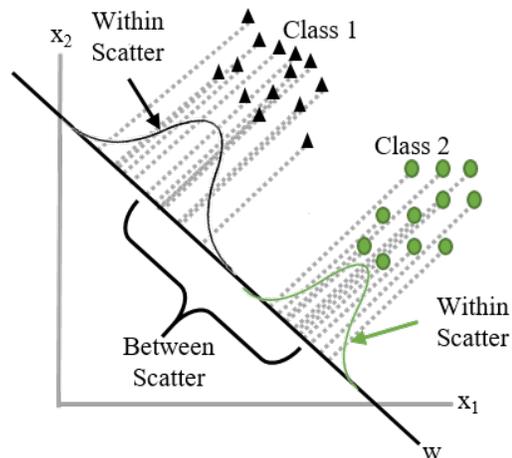


FIG. 2. Visual representation of the LDA technique.

they were not actively considered in the analysis as they provide limited predictive power given the size of the data-set relative to the number of available features. The academic performance features listed above are purposely selected to assist in determining if male and female students place differing importance on their academic performance when deciding whether to continue within CS. Additionally, the feature space provides consideration of the available pathways when deciding to change majors. For example, differing performance outcomes between CS courses and other technical courses ($cGPA > mGPA$) might prompt a student to change from CS to another STEM field. Similarly, should a student perform worse in their technical courses compared to other courses (as indicated by $GPA > fcGPA$) the student might consider leaving STEM majors all together. Providing for these different decision processes is intended to improve discrimination between persisting and non-persisting students.

We apply linear discriminant analysis (LDA) to the student data to extract a GPA-based decision model describing persistence in CS [9]. Like logistic regression, LDA is a supervised classification algorithm that makes use of weighted regressors on the feature space. In the case of LDA, weighting coefficients w_j are selected to maximize the ratio of the between-class separation and the within-class separation of two data classes when projected onto a subspace. Figure 2 provides a visual representation of the LDA process and outcome; ultimately the goal is to maximally separate the two classes (e.g. persist vs not persist) with the least amount of overlap. After applying LDA, each student profile can be transformed according to a scoring function of the form

$$\hat{x}_i = b + \sum_{j=1}^J w_j x_{i,j},$$

where w_j is the weighting on each descriptor $x_{i,j}$ (e.g. hs-GPA, GPA, mGPA, etc) and b represents a constant intercept adjustment. When used for predictive classification, if

weighted score for a student satisfies the inequality $\hat{x}_i > 0$, the student is predicted to persist. Conversely, for $\hat{x}_i < 0$, the classification model would predict the student to change majors. As such, the values w_j and the intercept b serve as weightings within a decision process model based on academic performance indicators.

The data set is constructed so that the LDA process enables the decision model to predict if a student will leave the CS major the following semester or if the student will ultimately go on to graduate. Accordingly, non-persisting student profiles are matched with the profile of graduating students; this requires the profiles of graduating students to be truncated to match the number of semesters enrolled of the non-persisting students.

III. RESULTS AND ANALYSIS

Figure 3 depicts the distributions of the weighted student scores following separate application of the LDA models for male and female students (with each green and red distribution conditioned according to actual persistence). Accordingly, two models are created: a male-specific model generated using male student data, and a female specific model generated using female student data. In each plot, the distributions of predicted outcomes for the male-specific and female-specific models are sub-divided according to actual outcomes (persist and not persist). The portion of the distributions to the right of the y-axis corresponds to student profiles predicted to persist; with the actual distribution of student profiles that did in fact persist and graduate in CS represented by the distributions above the x-axis. Circling a plot clockwise, each quadrant of the distributions finds correspondence to the True-Positive (upper-right quadrant), False-Positive, True-Negative, and False-Negative rates of the male-specific and female-specific prediction classifiers. In addition, Table III report the accuracy of each gender-specific model in predicting persistence outcomes. While the gender-specific prediction classifiers perform well at correctly predicting students who persist (94.7% and 99.5%), both the male-specific and female-specific models maintain high False-Positive rates; 30.3% of men and and 33.3% of women who did not actually persist were predicted to persist according to the gender-specific models. Such a result indicates that the GPA-based persistence models do not specifically describe or completely capture the decision process for a significant sub-population of students. That said, when compared with simpler models that only considers overall GPA and the full curriculum GPA (see accuracy results in Table IV), the full GPA-model (which uses many GPA features) provides a significant improvement in predictive power; major related GPA features appeared to have the greatest effect on prediction accuracy.

Direct comparison of the w_j and b coefficients for each gender-specific model is not particularly useful, as they represent the summation of weightings on individual compo-

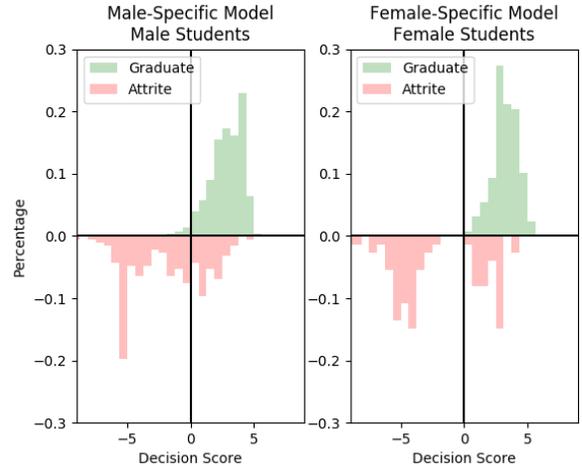


FIG. 3. Projection of male and female profiles according to gender-specific models.

TABLE III. Predictive Accuracy for full GPA Model.

Outcome	Male-Specific Model Prediction		Female-Specific Model Prediction	
	Persist	Not Persist	Persist	Not Persist
Persist	.947	.053	0.995	.005
Not Persist	.303	.696	.333	.666

nents (GPA, mGPA, hsGPA, etc.) and any interactions (e.g. GPA-mGPA, fcGPA-mGPA) that are also adjusted according to population means. However, the gender-specific decision models can be compared by using the one gender-specific model to predict outcomes of the other gender. The result of this procedure yields a high agreement rate between the gender-specific models. For 87% of the 465 male student profiles considered, the male-specific and female-specific models produced the same prediction. In all cases corresponding to the 13% difference, each profile is predicted to not persist in the male-model, while predicted to persist according to the female specific model. Similarly, for 84% of the 202 female student profiles, both models yielded the same persistence prediction. Here, the same bias occurs in 16% of the female profiles; the female-specific model predicts persistence, while the male-specific model does not. This bias in predictions does not indicate that male students have a higher threshold for persistence (i.e. men are more likely to change majors),

TABLE IV. Predictive accuracy for basic GPA model.

Outcome	Male-Specific Model Prediction		Female-Specific Model Prediction	
	Persist	Not Persist	Persist	Not Persist
Persist	.917	.082	1.000	.000
Not Persist	.389	.616	.407	.592

rather the disagreement is due to lower accuracy in predicting non-persistence for women, as indicated by the higher false-positive rate (see Table III). The high agreement rate in cross-application of the prediction models suggests that for the majority of students the decision to persist or not according to academic factors is not highly gendered.

IV. DISCUSSION

The results above provide meaningful opportunities for discussion. As stated, for the majority of male and female students, a decision process strictly based on GPA measures does not appear to be gendered as both male-specific and female-specific models make similar persistence predictions. We believe this result implies that the bulk of male and female students maintain similar scales of self-assessment in reference to persistence decisions within the CS major. Unlike [8], we did not find significant gender sensitivities to academic performance. That said, there remains significant limitations regarding the supporting data and the resulting models.

First and foremost, the GPA-based models over predict persistence, particularly for female students. This artifact is reflective of students with similar profiles who have differing persistence outcomes. This stands in contrast to an implicit modeling assumption for any classifier – which is certainly false – that students maintain the same threshold for accept-

able academic performance when deciding to continue or not within a major. The greater the real-world variation in this acceptable-performance threshold, the greater the confusion when predicting student outcomes.

With regards to data, the LDA modeling procedure excludes two factors that would likely improve the predictive power of the models. First the models do not consider potential bias generated through academic scholarships (like Florida’s Bright Future program) that requires students maintain a minimum GPA. Additionally, the models do not consider performance on UCF’s CS foundations exam. Computer science students are required to pass the exam within three attempts before they are formally declared to be CS majors and permitted to enroll in upper-division coursework. It is uncertain if male and female students respond equally to passing or not passing the foundations exam. Inclusion of these factors will likely improve the ability of the models to predict persistence according to academic measures.

Ultimately, the weakness in the predictive power for the female-specific model heavily hints at external factors driving the decision to persist in a major. Examples factors include but are not limited to: enjoyment, acceptance, and numerous gender specific issues. These factors were not included within the models, as they are not available. However, they do suggest opportunities for change to improve female retention in computer science.

-
- [1] National Science Foundation, National Center for Science and Engineering Statistics, *Women, Minorities, and Persons with Disabilities in Science and Engineering: 2017*, Tech. Rep. NSF 17-310 (Arlington, VA, 2017).
 - [2] Guili Zheng, Maguel A. Padilla, Timothy J. Anderson, and Matthew W. Ohland, “Gender Differences In Major Selection And Academic Success For Students Leaving Engineering,” in *Proceedings of the 2005 ASEE Annual Conference* (ASEE, 2005).
 - [3] Jennifer Wang, Hai Hong, Jason Ravitz, and Marielena Ivory, “Gender Differences in Factors Influencing Pursuit of Computer Science and Related Fields,” in *Proceedings of the 2015 ACM Conference on Innovation and Technology in Computer Science Education* (ACM, 2015) pp. 117–122.
 - [4] Adam V. Maltese and Robert H. Tai, “Pipeline persistence: Examining the association of educational experiences with earned degrees in STEM among U.S. students,” *Science Education* **95**, 877–907 (2011).
 - [5] Emily J. Shaw and Sandra Barbuti, “Patterns of Persistence in Intended College Major with a Focus on STEM Majors,” *NACADA Journal* **30**, 19–34 (2010).
 - [6] Ben Ost, “The role of peers and grades in determining major persistence in the sciences,” *Economics of Education Review* **29**, 923–934 (2010).
 - [7] Ann L. Owen, “Grades, Gender, and Encouragement: A Regression Discontinuity Analysis,” *The Journal of Economic Education* **41**, 217–234 (2010).
 - [8] Kevin Rask and Jill Tiefenthaler, “The role of grade sensitivity in explaining the gender imbalance in undergraduate economics,” *Economics of Education Review* **27**, 676–687 (2008).
 - [9] Maja Pohar, Mateja Blas, and Sandra Turk, “Comparison of logistic regression and linear discriminant analysis: a simulation study,” *Metodoloski zvezki* **1**, 143–161 (2004).