

Next-Term Student Performance Prediction: A Recommender Systems Approach*

Mack Sweeney
Computer Science
George Mason University
Fairfax, VA, USA
msweene2@gmu.edu

Jaime Lester
Higher Education
George Mason University
Fairfax, VA, USA
jlester2@gmu.edu

Huzefa Rangwala[†]
Computer Science
George Mason University
Fairfax, VA, USA
rangwala@cs.gmu.edu

Aditya Johri
Information Sciences
George Mason University
Fairfax, VA, USA
ajohri3@gmu.edu

ABSTRACT

An enduring issue in higher education is student retention to successful graduation. National statistics indicate that most higher education institutions have four-year degree completion rates around 50%, or just half of their student populations. While there are prediction models which illuminate what factors assist with college student success, interventions that support course selections on a semester-to-semester basis have yet to be deeply understood. To further this goal, we develop a system to predict students' grades in the courses they will enroll in during the next enrollment term by learning patterns from historical transcript data coupled with additional information about students, courses and the instructors teaching them.

We explore a variety of classic and state-of-the-art techniques which have proven effective for recommendation tasks in the e-commerce domain. In our experiments, Factorization Machines (FM), Random Forests (RF), and the Personalized Linear Multiple Regression model achieve the lowest prediction error. Application of a novel feature selection technique is key to the predictive success and interpretability of the FM. By comparing feature importance across populations and across models, we uncover strong connections between instructor characteristics and student performance. We also discover key differences between transfer and non-transfer students. Ultimately we find that a hybrid FM-RF method can be used to accurately predict grades for both

*(Does NOT produce the permission block, copyright information nor page numbering). For use with ACM_PROC_ARTICLE-SP.CLS. Supported by ACM.

[†]Corresponding Author.

new and returning students taking both new and existing courses. Application of these techniques holds promise for student degree planning, instructor interventions, and personalized advising, all of which could improve retention and academic performance.

Acknowledgments

This research was funded by NSF IIS grant 1447489.