



Predicting Student Outcomes to Drive Proactive Support:

An Exploration of Machine Learning to Advance Student Equity & Success

Prepared by

Nathan Greenstein, Grant Crider-
Phillips, Claire Matese, Sung-Woo Cho

Academic Data Analytics
Briefing Document
December 2021

EXECUTIVE SUMMARY

The University of Oregon's Academic Data Analytics Unit (ADA) is pleased to share preliminary results from a project that harnesses machine learning in service of student equity and success. In collaboration with the Division of Undergraduate Education and Student Success (UESS), ADA began with two ambitious goals. First, we sought to help UESS support first-year students proactively rather than reactively by predicting winter term retention before students matriculate. Second, we set out to pilot a process of transparent, responsible machine learning performed by the University's own experts. To date, this project – the first of its kind at the University – has achieved unprecedented progress towards both goals.

Using a process that prioritizes equity and trust, ADA developed a machine learning model that predicts winter term retention several times as successfully as alternatives, working especially well for potentially vulnerable groups. To ensure that our work purposefully tackles existing inequities instead of unintentionally cementing them, our team continually sought UESS's expertise on student and advisor needs to frame the project, review data, and evaluate model performance through an equity lens. Ultimately, ADA's model is two to four times as successful as alternatives, with performance still improving. All told, this project has shown that it is possible to leverage machine learning, put inclusivity and transparency first, and achieve meaningful results. Next, ADA looks forward to improving the winter term retention model, identifying other needs for predictive analytics at the University, and sharing the products of our work with peer institutions.

INTRODUCTION

Research shows that timely interventions can help institutions like the University of Oregon improve student success and achieve more equitable outcomes. At their best, such interventions are provided directly to the students who need them most, and are proactive rather than reactive (Astin, 1984; Kuh et al., 2011; Tinto, 1987). Taken together, these two points present a challenge: it is difficult to know which students are likely to struggle before they are struggling. This is especially true when considering student outcomes like retention, which are driven by a host of academic, cultural, and financial factors. To tackle this challenge, the Academic Data Analytics Unit (ADA) began a collaboration with the Division of Undergraduate Education and Student Success (UESS) in the fall of 2021. Together, we set out to use cutting-edge machine learning to predict which first-year students are at risk of not

returning for their second term, and to present our predictions *before* these students matriculate, allowing UESS ample time to offer support proactively.

This project has two key goals: to serve the student community by facilitating timely intervention, and to model a process of transparent, responsible machine learning performed by the University's own subject matter experts. Our machine learning model is designed to be used at the end of each summer to identify a relatively small number of students who are predicted least likely to return for their second academic term. During the first weeks of fall, UESS and other advisors can use this resource to focus their outreach such that students with higher predicted risk receive the support they need to remain enrolled through the winter and beyond. While developing the model, ADA worked closely with UESS to prioritize transparency and equity. We included a range of voices when making key decisions, and carefully tested our model for algorithmic bias, to ensure that our work seeks to actively remedy existing inequities instead of reinforcing them.

Although this project remains under active development, our initial results are positive. To drive our machine learning model, we identified 94 variables that stakeholders believed might be related to second-term retention. Drawing on ten years of historical data, our model uses these variables to predict retention two to four times as successfully as alternatives, with performance still improving. The remainder of this brief will offer more detailed discussion of the motivation and context for this project, our novel approach to the machine learning process, the initial results, and our conclusions and next steps.

MOTIVATION & INSTITUTIONAL CONTEXT

The literature tells us that early intervention can drive powerful improvements to student success and equity, but logistical realities force institutions to prioritize which students receive a given intervention earliest, or at all. Negative outcomes experienced at the beginning of a student's college career can disrupt their progress in ways that are difficult to recover from. By the same token, interventions tend to be most impactful when they focus on preventing a negative outcome, rather than reacting to it (Astin, 1984; Kuh et al., 2011; Tinto, 1987). However, constraints on resources like staff, facility space, and funding mean that not every student can receive every intervention, and not all students can receive any given intervention at the same moment in time. These facts make it important to have timely insight into which students are at greatest risk for negative outcomes, especially since the

burden of such outcomes tends to fall disproportionately on the most vulnerable students, compounding existing inequities (Swail et al., 2003).

Non-retention is damaging to both individual students and the University as a whole. In the case of this project, UESS strives to prevent negative outcomes by offering advising services to first-year students very early during fall term. However, because there are naturally far fewer advisors than students, not all students can meet with their advisors on this accelerated timetable. To have the greatest possible impact, UESS seeks to identify students are at highest risk for negative outcomes and offer them prioritized access to advising.

Before this project, UESS prioritized students with low predicted first-year GPAs, but viewed winter retention as a more important outcome. The existing process relied on a linear regression model that sought to predict first-year college GPA based on eleven variables, with high school GPA the most influential among them. Although UESS hoped to instead prioritize their early advising intervention based on winter retention likelihood, the existing model did not predict retention accurately enough to allow this. This limitation is likely driven by several factors. Most simply, it is rare for students to not return for the winter – only approximately 4% did not return over the past ten years. Moreover, retention is a relatively erratic outcome, as it is driven by a blend of academic, social, cultural, familial, and financial factors unique to each individual student (Ziskin et al., 2009). These difficulties are compounded by the fact that predictions must be created before students matriculate, when only limited data is available, in order to support the accelerated timeline that is essential to this intervention.

This challenge – to predict a mercurial outcome before any data on students' college careers exists – struck ADA as one that might only be met with the help of machine learning. However, some stakeholders maintain healthy skepticism towards machine learning, especially with respect to transparency and equity. Such concerns are certainly understandable given the past shortcomings of high-profile projects by technology companies, but we believe that such issues are consequences of faulty processes, not inherent to machine learning. Like many tools, machine learning can do good or harm, depending on how it is wielded. We set out to wield it for good, meeting a need for challenging predictions, and creating a process that is transparent, inclusive, and fair.

A PROCESS FOR RESPONSIBLE MACHINE LEARNING

Past applications of machine learning, both within the University and beyond, can arouse reasonable misgivings among those concerned with accountability and student equity. One such example is a machine learning solution for predicting student risk that was produced for the University by a third-party vendor. Due to its proprietary nature, this product offered sparse visibility into how it was developed, which variables it considered, and how well it performed for students belonging to different groups. The lack of transparency into the model's development led to a lack of trust in the model's ability to support equitable advising practice, which undermined confidence in the model. More publicly, it has come to light that certain machine learning solutions serve some groups better than others (Hardesty, 2018), or worse, exacerbate systemic inequities (Angwin et al., 2016; Dastin, 2018). Those aware of these issues might reasonably worry about the unintended consequences of a new machine learning project. We applaud this regard for the wellbeing of potentially vulnerable students, and we set out to complete this project in a way that puts concerns like these front and center.

From day one, active collaboration with UESS has shaped ADA's predictive modeling process, demonstrating a key advantage of performing such work in-house. A machine learning solution is most valuable when designed around the needs of those who will use it, and we believe that it is only truly successful if it inspires trust in users and other stakeholders – in this case, advisors and students. To that end, ADA and UESS convene frequently, and each group openly asks questions and offers feedback to the other. As decision points arise, ADA outlines technical considerations and tradeoffs, and UESS responds based on their knowledge of student and advisor needs. This process plays to both teams' strengths, and it has allowed us to answer fundamental questions like the following (plus myriad smaller ones) with full transparency and the trust that it fosters:

- What limitations of the existing system should we seek to improve upon?
- How early should predictions be available each year?
- What is the maximum number of at-risk students who should be identified?
- What information about each student should be provided to advisors?
- Which predictor variables can we responsibly include in our model?
- How should we define equity in the context of this model?
- How should we measure and report on this model's predictive performance?

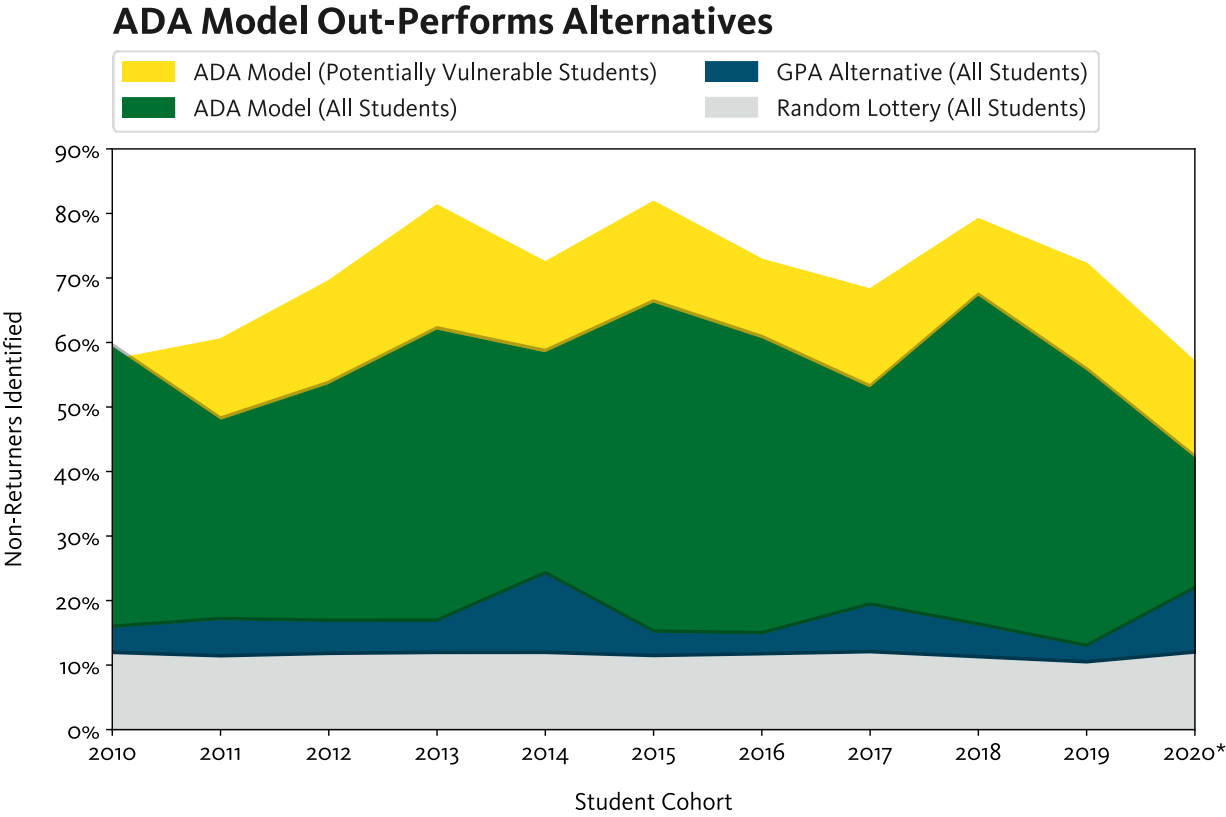
This inclusive and transparent approach forms the bedrock of our process, and we continue to apply it through each stage of our model's development. To gather data, we worked with numerous offices within the University, including Admissions, Financial Aid, Information Services, Housing, Clark Honors College, Student Orientation Programs, First Year Programs, Athletics, and the Registrar. In addition, we assembled publicly available data to enrich the model, with a particular eye for capturing social disparities such as geographic differences in upward mobility. Throughout, UESS has reviewed the data we select to ensure that it is suitable from a student equity perspective and will not raise concerns or undermine the trust of students and advisors. Finally, as our model has begun to produce results, ADA and UESS have worked together to ensure that it is fair and equitable. ADA has outlined possible theoretical approaches to defining and measuring algorithmic bias, and UESS has helped select the approaches most relevant to this particular context, including potentially vulnerable groups to consider. ADA has then evaluated the model to ensure that it meets or exceeded UESS's criteria for fairness and equity.

INITIAL RESULTS

Initial results from ADA's machine learning model are positive, and we expect performance to continue improving. Currently, the model successfully predicts winter term retention two to four times as well as alternatives, and it performs especially well for potentially vulnerable students. The model leverages 61 variables from University data and 33 variables from public data, captured over the ten years represented by the 2010 through 2019 student cohorts. Like all machine learning models, ours mines historical data for mathematical relationships between predictor variables and the outcome of interest, then "learns" these relationships so that it can predict future outcomes when only the predictor variables are known. Internally, the model is built on the XGBoost machine learning framework, which seeks to produce accurate and generalizable predictions by combining many simple models into a single "ensemble" model that is designed to perform better than any one model could on its own.

To evaluate the model's performance, we compare its predictions to actual winter term retention outcomes over the same ten-year period. Additionally, we apply the same comparison between predicted and actual outcomes within the 2020 student cohort, which was withheld from the model during development. This step is designed to simulate the model's intended use for incoming student cohorts, which is typically a more challenging prediction task. In discussing the model's performance,

we present the proportion of students who did not return for winter term (“non-returners”) that the model identified in its list of students at greatest risk of not returning. We then contextualize these results by comparing them to the same metric as taken from two alternative approaches, one based on a random lottery, and one based on high school GPA.



* The 2020 cohort was hidden from the model during development to simulate an incoming cohort of students.

For cohorts 2010 through 2019, ADA’s model identifies an average of 58% of non-returners, with individual years ranging from 48% to 67%. For the 2020 cohort, ADA’s model identifies 42% of non-returners. Focusing specifically on potentially vulnerable students – defined here as first-generation students and those belonging to traditionally underserved races and ethnicities (“URMs”) – the model identifies an average of 71% of non-returners in 2010 – 2019, and 57% of non-returners in 2020. For comparison, a hypothetical approach based on a random lottery would identify only 11% of non-returners, and a hypothetical approach based directly on high school GPA would identify only 17% of

non-returners.¹ In other words, by identifying students using ADA's model instead of identifying students with low high school GPAs, early results suggest that UESS could proactively reach out to between 2.0 and 4.5 times as many non-returners overall, and between 2.4 and 4.6 times as many potentially vulnerable non-returners.²

We are pleased with these initial results, and we plan to continue increasing the model's performance. Even a twofold improvement over the GPA alternative has potential to make a deeply meaningful difference in students' lives, but we believe we can go further. We continue to actively search out additional data and refine the modeling algorithm's development, with a particular emphasis on predictions for new student cohorts like 2020. We look forward to sharing updates on improved model performance over time, as well as diving deeper into modeling challenges such as uncertainty created by the COVID-19 pandemic.

CONCLUSIONS & NEXT STEPS

The Academic Data Analytics unit began this project with two goals: to advance student success and equity by predicting a complex outcome, and to pilot a novel approach to responsible, in-house machine learning. Although work remains to be done, we have achieved unprecedented progress towards both goals. Our model's predictions of winter term retention are two to four times as successful as alternatives, and performance is still improving. Moreover, to build the model, we developed a machine learning process that is radically different from what the University had experienced before. All told, this project has shown that it is possible to leverage machine learning, put inclusivity and transparency first, and achieve meaningful results even when available data is limited. Our team is excited to continue refining and building upon this work.

Looking ahead, ADA plans to focus on improving the winter term retention model, identifying other needs for predictive analytics at the University, and sharing the products of our work with our peers. To

¹ For reasons that we have not yet explored, high school GPA appears to be more accurate than usual among the 2020 cohort. This may be related to the COVID-19 pandemic.

² These initial results have not yet been subject to the statistical testing necessary to supply detailed estimates of such outcomes in a range of scenarios; these values are observed directly from our model's current performance.

date, our team continues to search actively for additional data to enrich the model developed under this project, and to implement algorithmic improvements to draw better performance from existing data. More importantly, we continue to collaborate with UESS to expand our evaluation of how well the model serves potentially vulnerable groups, and we are prepared to make any corrections that become necessary. We understand that first-generation and URM identities are by no means the only facets of potential vulnerability, and we are committed to ensuring that our model advances equity in a holistic sense. Should we discover that the model underserves a potentially vulnerable group, we will correct this by refining the modeling process to place additional emphasis on data from members of such groups, re-balancing its performance.

Meanwhile, ADA leadership is reaching across University offices to identify the next challenges which might best be met with machine learning, and we look forward to publicizing takeaways from our work. We aim to share both technical and procedural findings so that peers can replicate and build off of this project's achievements. We hope that, by conducting and sharing the results of this work, we can help drive cutting-edge gains in student success and equity at a variety of institutions, regardless of their capacity to carry out machine learning from the ground up.

WORKS CITED

- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). Machine Bias. *ProPublica*.
<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing?token=p-v0T1xjfOJ8jrJHzc08UxDKSQrKgWJk>
- Astin, A. W. (1984). Student Involvement: A Developmental Theory for Higher Education. *Journal of College Student Personnel*, 25(4), 297–308.
- Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- Hardesty, L. (2018, February 11). Study finds gender and skin-type bias in commercial artificial-intelligence systems. *MIT News | Massachusetts Institute of Technology*.
<https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212>
- Kuh, G. D., Kinzie, J., Schuh, J. H., & Whitt, E. J. (2011). *Student Success in College: Creating Conditions That Matter*. John Wiley & Sons.
- Swail, W. S., Redd, K. E., & Perna, L. W. (2003). *Retaining Minority Students in Higher Education: A Framework for Success: ASHE-ERIC Higher Education Report*. Wiley.
- Tinto, V. (1987). *Leaving College: Rethinking the Causes and Cures of Student Attrition*. University of Chicago Press, 5801 S.
- Ziskin, M., Hossler, D., & Kim, S. (2009). The Study of Institutional Practices Related to Student Persistence. *Journal of College Student Retention: Research, Theory & Practice*, 11(1), 101–121.
<https://doi.org/10.2190/CS.11.1.f>