Artificial Intelligence and Student Success

Help Every Student Stay on Track

Artificial intelligence (AI) is revolutionizing computer science—and the world around us. Its potential to sort through vast amounts of data —and learn to make intelligent interventions from that data —is something that we now take for granted every time a recommendation engine serves up a new book for us to read or a new song for us to download.

AI, as we explain in this white paper, has the potential to help us redefine how we support student success —by informing policy, identifying success strategies for at-risk groups, and helping *every student* stay on track. More than the engine behind the ubiquitous "virtual assistant," today's AI tools have the potential to turn student success tactics into personal and powerful interventions—ones that reach each individual student where they are to make immediate and lasting impacts.

AI has the potential to turn student success tactics into personal and powerful interventions—ones that reach students where they are to make immediate and lasting impacts.

Student success is a complex subject. How we define it, how we measure it, and, most important, how we achieve it, has occupied educators, policy makers, and advocates for more than a decade. Historically, student success was all about degree completion, but soon we saw performance funding tied to more granular indicators like retention, persistence, and time-to-degree.[1] More recently, transfer rates from community colleges to four-year institutions[2] and career readiness measures have caught the attention of researchers.[3] Today, we see a new call for a more holistic approach to student success—one that considers the experiences of increasingly diverse learners whose education goals are as varied as the barriers they encounter when trying to achieve them.[4]

As ideas about student success evolved, strategies shifted. We are more aware that first-generation students, underserved students, and working adults require different resources to succeed. At many institutions a focus on programming for first-year students has expanded to encompass the full student experience. Educators are generally more willing to engage with new technologies to deliver better support by tracking degree progress, providing alerts to faculty and staff, and assisting student advisors with their work. While these solutions are an important part of the student success ecosystem, technology can offer more.

[1] https://www.thirdway.org/report/lessons-learned-a-case-study-of-performance-funding-in-higher-education
[2] https://www.jkcf.org/research/persistence/
[3] https://ccrscenter.org/
[4] https://er.educause.edu/articles/2019/10/student-success--3-big-questions

Supporting students through their education journeys has been challenging in large part because educators don't have the tools they need to make sense of the myriad factors that influence a student's likelihood to succeed. While most institutions have stepped up to implement student success programs and policies, those efforts still often fail to address the needs of *particular students* who are struggling with their own particular set of issues. As the conversation about student success continues to mature, advocates and educators are asking institutions to look beyond categories and frameworks that may cause us to "inadvertently…overlook the circumstances of students' lives" and seek better ways to address individual student needs.[5]

In 2020, the global pandemic is disrupting students' lives in profound ways. How they will navigate their own personal student journeys is yet to be written—but they will likely need the assistance of a more robust and nimble support network. AI-enabled solutions can help institutions predict the acute pains that will arise and mitigate their impact on each student in a personal and empowering way. Today's AI can help every institution help students find a better way forward.

Student success efforts often fail to address the needs of particular students who are struggling with their particular set of issues.

Why AI?

Most of us are wired with brains that can assess a small number of variables, identify how those variables have produced a particular problem, and identify a solution to that problem. We see a crying child in one aisle of the supermarket and a distraught woman in the next. Problem: lost child. Solution: point mother to child. The relationships between these variables tend be linear and we can identify cause, effect, and solution quickly.

By the same token, we're good at understanding the linear correlation between grades and dropout rates. We use that metric as an indicator to intervene with a student whose 2.0 GPA raises a red flag. We're successful—at least in part. Institutions have applied the logic of thresholds across their entire student body in order to effect change. It's one of the few ways they have to work at scale to support as many students as they possibly can.

How do we move from thinking about student success in the aggregate to a focus on a particular student's needs?

But the variables that affect a particular student's educational success are often deeply complex, involving hundreds of data points. While thresholds are useful for helping institutions intervene somewhat successfully with a cohort of students, they have had more difficulty intervening very successfully with every *individual* student. If institutions could identify those variables, they might discover that a student whose grades are well above that 2.0 GPA threshold is facing quite different hurdles: her home might be 300 miles from campus in a rural part of the state, her unmet need might be much higher, and she may not have had the means to purchase a meal plan. While her GPA is 3.2, her likelihood of dropping out—for reasons other than academic ones—could well be greater than the student with a 2.0 GPA.

[5] http://download.hlcommission.org/initiatives/StudentSuccessConversation.pdf

How, then, do we deliver the right interventions at the right time to *every student* to ensure their success? How do we move from thinking about student success in the aggregate to a focus on a particular student's needs? The scalable pathway to answering these questions is AI-driven analytics. As computing power and data storage become more accessible to institutions of every kind, AI provides a new way to process, understand, and act on the very complex variables that contribute to student success.

PREDICTIVE ANALYTICS

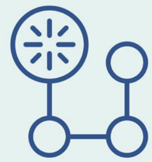

PRESCRIPTIVE ANALYTICS

EXPLAINABLE AI

Predictive, Prescriptive, and Explainable

While we've seen AI applied in discrete and limited ways at different points across the student journey, AI has the potential to do much more. First a quick tutorial on a few important terms:

Predictive Analytics

This term describes machine learning techniques that uncover features and patterns in data in order to make predictions based on those patterns. Predictive analytics learns from the data it ingests rather than performing operations scripted by a program. It can sort through vast amounts of data and combine that behavioral history in new ways to identify the variables that predict success. For example, predictive analytics might look at data indicating a student is only visiting the dining hall every third day and combine that data with other variables, such as that student's unmet need. By doing so it identifies the risk facing that student: food insecurity. Predictive analytics rapidly sorts through variables to identify future risks in a way the human brain can not.

Prescriptive Analytics

Predicting who is most at risk is only half of the equation. Educators want to know what actions to take to increase the probability of that student's success. Prescriptive analytics sorts through the many interventions that will have an impact on student success to identify those which will have the greatest impact. While predictive analytics can tell you why a student is at risk, prescriptive analytics tells you what you can do to mitigate those risks and deliver the best possible outcome.

Explainable AI

In the business world, AI is often used to analyze, predict, and influence buyer behavior. Few businesses have a mandate to understand how an algorithm comes to its conclusions. In the world of higher education, there are a host of reasons for needing a more complete understanding of how these conclusions are reached. Explainable AI (as opposed to 'black box' AI) exposes its rules so that they can be understood, explained, and audited.
(Read more about explainable AI here.)

Institutions can now make student success policies truly actionable - for each student.

At Othot, we think the combination of these three AI attributes delivers powerful new ways to make intelligent interventions in every student's journey. Institutions now have the ability to make their student success policies truly actionable—for every student. By applying predictive and prescriptive analytics *at the individual level*, your institution is not simply analyzing the risk factors most likely to derail a cohort of students, but those most likely to derail Jacob or Martina or Kyla or Xavier. And you are not only creating institution-wide success programs (as important as those are), but you are designing *specific interventions* that are most appropriate for Jacob and Martina and Kyla and Xavier. And finally, by making those operations transparent and explainable, your institution has better tools to build and maintain the level of trust that your stakeholders expect.

AI and Student Success: Two Views

AI is very good at surfacing variables that have a significant impact on student success. But it's even better at helping us see how those variables can be combined in surprising ways to identify previously unknown risks. In the two views below, we've used sample data to illustrate how explainable AI can help you understand persistence at a hypothetical school with 1,500 students. Both views illustrate AI's predictive capabilities. Both are also examples of how explainable AI details the variables used in those predictions as well as the relative weight of their impacts.

The Aggregate View

The Aggregate Impacts

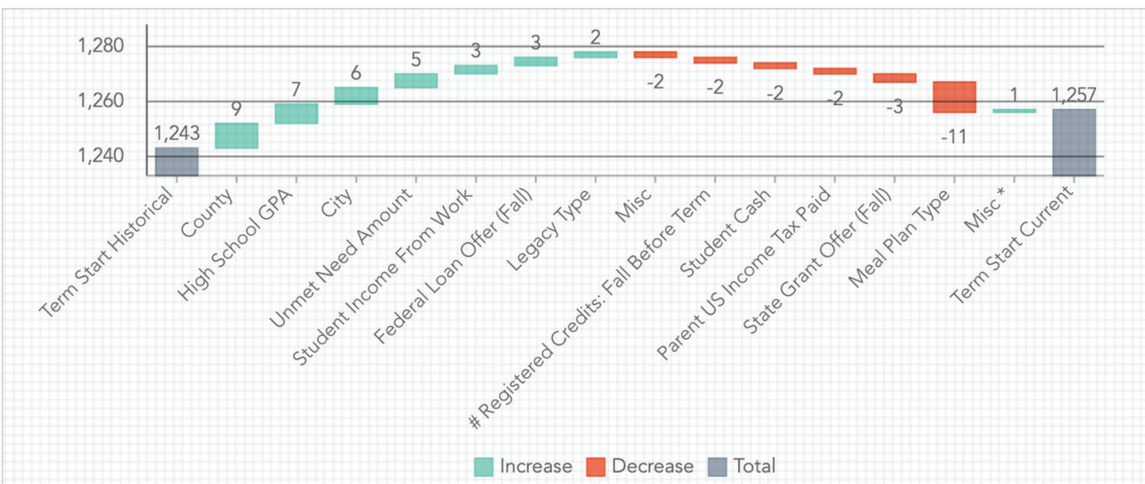

Figure 1 - Waterfall of the Student Body

The waterfall table in Figure 1 summarizes the aggregate reasons that students at this institution are failing to persist. The gray bar on the left shows the historical yield of the student body that persisted to the next fall semester. The gray bar to the right shows the predicted yield of the student body likely to persist to the next fall semester. The bars in between show the positive (turquoise) and negative (red) impacts of the top variables impacting this student cohort.

We can see that academic preparedness ("High School GPA") has a significant positive impact, accounting for a hefty 7-point increase in likelihood to succeed. But we see a different story when we examine the red bars on the right, which expose a series of variables whose negative impacts will potentially erase what we often assume is the best predictor of success. Most of these variables are associated with unmet need or food insecurity. This provides the institution with the insight it needs to affect changes to policy that will have a material impact on increasing overall retention and persistence.

This view of the aggregated data is essential for understanding how and why students across your institution are or are not succeeding and what kinds of interventions will be the most successful. The aggregate view of these kinds of data can help inform policy, design roadmaps for refining those programs, and meet increasingly rigorous expectations for accreditation and funding.

The Individual View

As informative as this aggregate view is for institutions seeking broad solutions to student persistence, AI can also be used to examine why a particular student does or does not persist. Figure 2 demonstrates how AI can surface the combination of variables that can potentially derail an individual student.

Individual Score Impacts

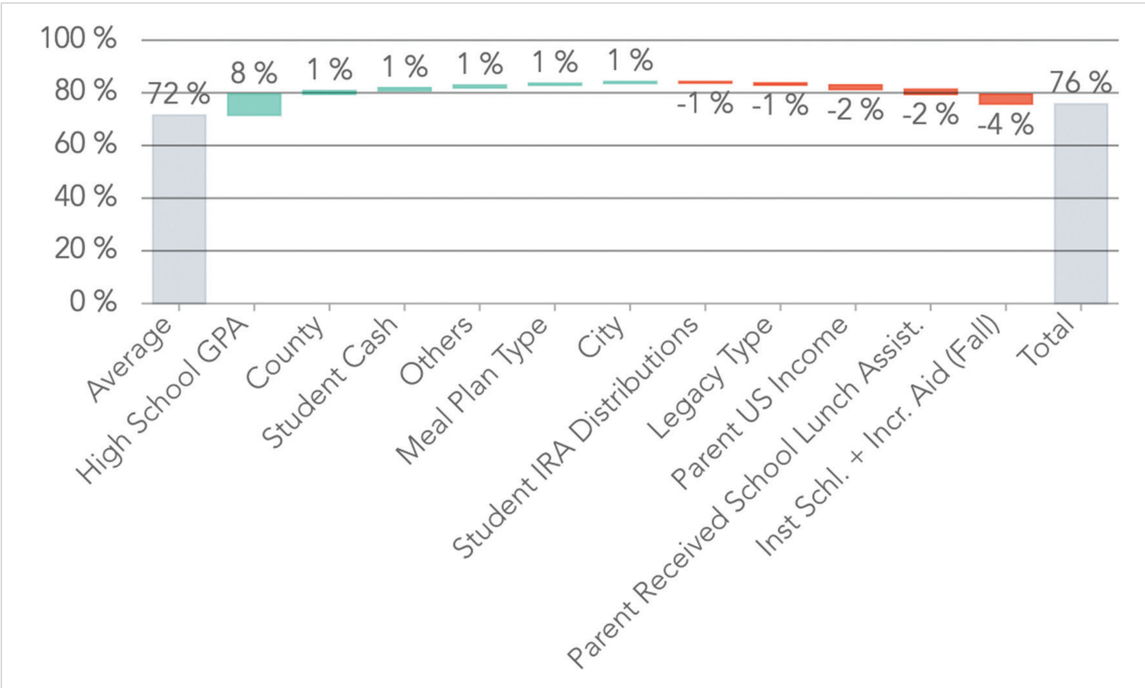

Figure 2 - Waterfall of an Individual Student

The left and right gray bars indicate the institution's average retention rate and this specific individual's retention rate, respectively. As in our example above, GPA is weighted heavily—the better a student's GPA, the more likely he or she is to persist. But, while this student's likely outcome is slightly better than average, risks still loom. And the model tells us why: of the five negative variables in red, four are related to financial stability.

AI can be used to examine the reasons why a particular student does or does not persist.

This view of this particular student's risk factors underscores the very real impact that unmet need is having on his or her chances of success. By improving this student's financial situation, the institution can increase those chances by nine points (i.e., the sum of the negative variables related to financial need). Because the factors having the most impact are all related to money, the institution can align the right kind of support much more quickly—by providing emergency funds or redesigning next semester's financial aid package, for example. As well, with this data and its impact more visible, the institution can further refine its aid policies to increase every student's likelihood to persist.

Improving Outcomes: From Predictive to Prescriptive Analytics

In the examples above, we see how AI can help predict student success risk factors at both the aggregate and individual levels. As institutions look for new ways to help more students succeed, that data will be critical. But predictive analytics, as useful as it is in identifying problems, doesn't tell us *how* to solve those problems. Prescriptive analytics does.

Prescriptive models governed by AI can consider significantly more data faster and more comprehensively than traditional what-if scenario planning. And because AI "learns" from the data it consumes, it can draw inferences from even incomplete data sets, delivering more precise and relevant recommendations.

Predictive and prescriptive analytics help you understand what happened, what is happening, what is going to happen, and how to change it.

Student Likelihood to Persist Before Advisor Meeting

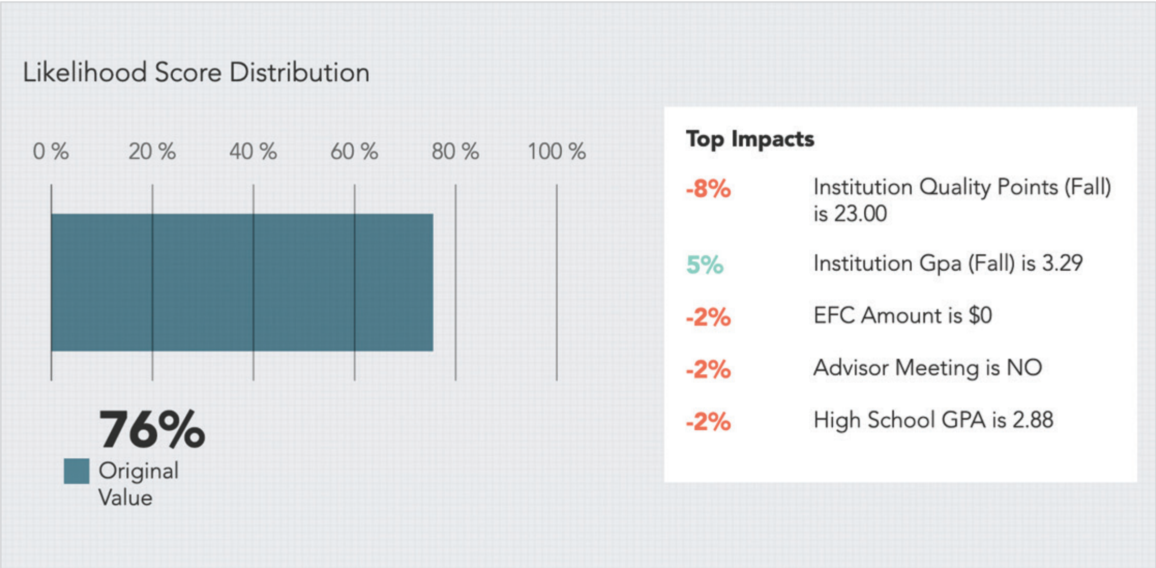

Figure 3 - Student Likelihood to Persist Before Advisor Meeting

In Figure 3, we see that the prediction for a student's likelihood to persist is 76%, a middling score produced by a number of different variables. We can also see that at least two negative variables (EFC Amount and Advisor Meeting) are actionable.

Student Likelihood to Persist After Advisor Meeting

A prescriptive AI model can deliver a precise recommendation for action.

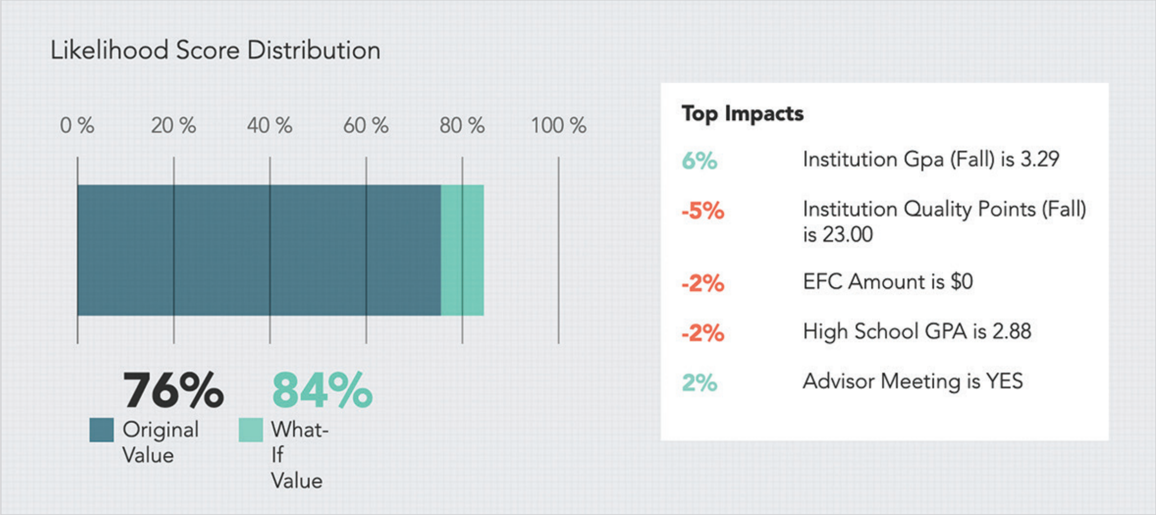

Figure 4 - Student Likelihood to Persist After Advisor Meeting

In Figure 4, we see that a prescriptive AI model can deliver a precise recommendation for action that will produce a measurable outcome. In this case, by providing this student with access to an advisor, her likelihood to persist increases by 8%, increasing the student's chances of success from fair to good.

With AI-driven solutions like these, we can understand the relationships between hundreds and even thousands of variables to understand and intervene in student behavior. And we can do that not just at one particular point in time, but in real time, across the student experience. As data about student interactions changes, AI-generated predictions and prescriptions will continue to learn and evolve, providing even better insight and smarter recommendations.

One of AI's most human features is its ability to learn and grow, and the sooner you start training your AI models, the more valuable they will become.

The Data Question: How much is enough?

Getting started with an AI project requires data—that's a fact. But building predictive and prescriptive AI models that will make a difference at your institution shouldn't take years. Most institutions have more data than they realize and, in our experience, can implement an AI platform like Othot's in about three months.

Institutions have been collecting data routinely to address reporting and accreditation requirements, improve enrollment, and provide better services to their education communities. External data sources are also valuable for feeding AI models; some of those sources are public, some are available from third parties. And of course, partners like Othot can also "engineer" your existing data to produce new variables from existing data. The point is: data is everywhere. It's simply of matter of understanding what data you need to begin.

There are three general categories of data that help inform AI models, all of which most institutions have readily at hand:

Demographic and personal data: Demographic and personal data: By the time a student is admitted to your institution, you've already gathered quite a lot of data, including home address, academic performance, standardized test scores, diversity data, programs of interest, unmet need, and more. Once a student enrolls, you gather information about on-campus life: housing, classes, grades, co-curricular activities, and more. Most of this information is readily available in an institution's SIS, CRM, and LMS.

Behavioral data: You also have an extraordinary amount of data about the actions your students are taking: the things they do and why they do them. You know, for example, how long it typically takes for a student to register for classes. That data becomes useful to AI models when that "typical" pattern changes, when a student who used to register for classes immediately registers a day before the deadline. Surfacing behavioral data sometimes requires thinking about data in a different way, such as recording frequency, intervals, and duration as in our example here.

Interventional Data: Students engage with your institution in a number of different ways. You likely reach out multiple times in a given school year through your CRM and at least some of those messages require a response. Whether a student responds or doesn't is a great data point for AI. So are any of the actions a student might take during the academic year: signing up for study abroad, registering for a review session, opting out of a meal plan. All of this data provides new information that AI models can learn from in order to predict outcomes and prescribe solutions.

Scale Your Student Success Efforts with AI

Even if you only have some of those data points at hand, AI can help. So too, can Othot. We've identified a core set of data points that most institutions can gather relatively easily for an initial AI project. We also bring external data sources to the engagement and combine those data points in different ways to create new variables. That means you will be able to make predictions and prescribe solutions with a modest amount of data and see progress quite quickly.

Over time, as you provide more data to the AI model, its predictions and prescriptions will become smarter, and the changes your institution can affect will continue to grow in impact. The point is to start early. Because one of AI's most human features is its ability to learn and grow. And the sooner you start training your AI models, the more valuable they will become.

If you would like to learn more about how you can improve student success at your institution and make a difference in the lives of every student, please contact us today at othotteam@othot.com or visit our website at othot.com.