CASE STUDY E: Transferring predictive models to other institutions from Marist College

Supported by a grant from EDUCAUSE's Next Generation Learning Challenges initiative, with funding from the Bill and Melinda Gates Foundation, the Open Academic Analytics Initiative (OAAI) sought to address three research questions:

1. What are the potential challenges, solutions, and benefits associated with developing a completely open-source early alert solution for higher education?
2. To what degree can predictive models be imported from the academic context (e.g., a four-year private liberal arts college) in which they were developed to new and potentially very different academic contexts (e.g., two-year community colleges)?
3. What intervention strategies are most effective in helping academically at-risk students succeed?

(Jayaprakash et al, 2014)

Building on open source software, this project has pioneered an open approach to learning analytics, and some of its tools are now being deployed as part of Jisc's learning analytics architecture.

Key takeaway points

» The most important predictor of future academic success was found to be partial contributions to the final grade

» The predictive models developed at one institution can be transferred to very different institutions while retaining most of their predictive abilities

» There are no gains from providing a sophisticated support environment for at risk students; simply making them aware that they are at risk may suffice

» In one course there was a significant improvement in final grade (6%) with those at risk students who were subject to an intervention compared with the control group who were not

» Interventions with students deemed to be at risk may have unintended consequences such as encouraging students to withdraw who might ultimately not have failed

Developing the predictive model

The initial predictive models were based on data from Marist College, a liberal arts institution in New York State. Building on Purdue's approach, data sources included demographic details such gender and age, aptitude data such as high school scores, and various aspects of VLE usage. These models were then deployed at two community colleges and two state universities with high numbers of ethnic minority students with low retention rates to investigate how portable the models were and whether the interventions were as effective. The models were then released openly with the intention of others taking them further. Jisc is now adapting the models for UK higher and further education with some contextual and linguistic changes.

As well as developing and transferring the predictive models, the other main aspect of the initiative was to investigate the effectiveness of the resulting interventions with at risk students. An Online Academic Support Environment was developed containing study skills materials and community support from paid specialists and student mentors.

Data used included the number of VLE course sessions opened by the student, the number of forum threads read and contributed, and the number of assignments and tests submitted. Gradebook scores entered by the instructor were combined to produce an aggregated score. The data was reviewed by IT staff to ensure its integrity. Attributes where data was missing in 20% of cases were discarded. Meanwhile the variability in usage of the VLE by different instructors was addressed with the decision that, for a VLE tool to be included in the data, at least 50% of the students should have used it at least once. It was recognised that the variability of workload, assessment and other activity across instructors and courses could have an impact on the transferability of the predictive models. This was addressed by replacing frequencies of VLE events by ratios and proportions, normalising them by dividing them by the average course frequency. Training the predictive model with historical data also involved removing any record of any student record there was no final grade recorded and removing "outliers", defined as "an observation distant 3+ standard deviations from the mean".

Sandeep Jayaprakash and his colleagues at Marist College deployed logistic regression to calculate whether students were either academically 'in good standing' or 'at risk'. They compared the resulting predictions with those produced from three other methods and found logistic regression appeared to be the best, more accurately predicting whether a student was at risk than the others. The predictive model itself is encoded in an XML-based language known as the Predictive Model Markup Language (PMML). This is the main standard available for storing and exchanging predictive models.

They discovered that the three most significant metrics for predicting student success were *marks on the course so far*, *grade point average* and *current academic standing*.

The model was trained with data from Marist College and rolled out to the partner colleges and universities. It was found to transfer well, with around 75% of at risk students being identified in three out of the four institutions. The researchers expected a much larger difference in how the institutions compared with Marist, given the difference in student cohorts, type of institution and teaching practices. They believe that theirs was the first project to identify *partial contributions to the student's final grade*, i.e. the results of assessments and assignments taken during the module and entered in the VLE gradebook tool, as the key indicator for determining at risk students. This enables instructors to take action much earlier in the semester than has been done previously. However they acknowledge that where partial contributions to the grade and cumulative GPA are not available in an institution the model way not be appropriate.

Interventions

Reports were produced showing the students deemed to be at risk, who were then subjected to one of two different intervention strategies. The first group was sent a message stating that they were at risk of not completing the course, and providing guidance on how they could improve their chances. The second was directed to the Online Academic Support Environment where open educational resources were provided in areas such as study skills, time management and stress reduction, as well as subject specific content in algebra, statistics, writing and researching. There were also mentoring opportunities from peers and support staff. Both interventions involve standardised text which becomes more serious in tone with every message.

Instructors had the say over whether these messages were sent to students who were identified as at risk at three points during the semester (25%, 50% and 75% of the way through it). Messages such as the following were available:

> *"Based on your performance on recent graded assignments and exams, as well as other factors that tend to predict academic success, I am becoming worried about your ability to complete this class successfully."*

and

> *"I am reaching out to offer some assistance and to encourage you to consider taking steps to improve your performance. Doing so early in the semester will increase the likelihood of you successfully completing the class and avoid negatively impacting your academic standing."*

Students were invited to visit their instructors during office hours, set up appointments with tutors or academic support staff, to consider participating in a study group, to view online materials, and to take additional exercises or practice tests.

Two different cohorts of 1,739 and 696 students in 2012 were divided into three groups: students who were sent a message, those who were directed to the support environment and a control group which received no interventions. No differences were found between the two treatment groups overall however there was a 6% improvement in final grade for those subjected to an intervention over the control group in one course.

Another finding was that withdrawal rates were larger in the intervention group than among control subjects. 25.6% of interventions students withdrew in contrast with only 14.1% of control students. However the results vary considerably between semesters and the authors refer to research at Purdue which showed similar inconsistencies. What differences there are may be explained by students preferring to withdraw early rather than running the risk of failing later. It is possible that the model will falsely identify students as at risk who might not have failed but are prompted to withdraw by being told they are at risk. The authors argue that this is a reason to strive to develop more accurate predictive models.

Conclusions

The Marist researchers conclude that the predictive model helps to provide students with earlier feedback on their progress, allowing them to address any issues before it is too late. While they found that many instructors welcomed the extra information on at risk students, it was not clear whether they would have been able to identify them anyway without the analytics. They also note that there are some students who appear to be "immune" to interventions. They found that very few students who failed to respond to the first intervention improved after the second or third.

References

Jayaprakash, S. M., Moody, E. W., Lauria, E. J. M., Regan, J. R., & Baron, J. D. (2014) Early Alert of Academically At-Risk Students: An Open Source Analytics Initiative. Journal of Learning Analytics, 1, 6–47. https://epress.lib.uts.edu.au/journals/index.php/JLA/article/view/3249