The Student Advice Recommender Agent: SARA

Jim Greer, Stephanie Frost, Ryan Banow, Craig Thompson, Sara Kuleza, Ken Wilson, and Gina Koehn

University of Saskatchewan, Saskatoon, Canada
{firstname}.{lastname}@usask.ca

Abstract: SARA, the Student Advice Recommender Agent is a system somewhat like an early alert system, where predictive models of learners' success combined with incremental data on learners' activity in a course can be used to identify students in academic distress. With SARA, rather than give alerts to academic advisors or professors, we provide personalized advice directly to students. An advice string – "A note from SARA" is prepared for each student every week in a semester-long course. The system attempts to direct students to appropriate learning supports and resources according to their individual needs. We have observed a significant year over year improvement in unadjusted student grades after the SARA's advice recommender was implemented in a 1200-student freshman STEM course.

Keywords: early alert, personalized advice, persona, recommender agent

1 Introduction

Early alert systems for students at academic risk have been in use for several years. In such systems, students who seem to be struggling in a course, as evidenced by lower term grades, minimal engagement in learning management system (LMS) activity, or low attendance may be issued warnings or alerts [1]. In most systems, instructors are involved in directing the delivery of alert messages. In some systems, these alerts are also issued to academic advisors (as in Starfish Early Alert or Ellucian Student Success) so that follow up appointments with an advisor or learning specialist can be booked if the advisor so wishes. For the most part, students who seem to be minimally engaged or who are falling behind in coursework, or who are failing intra-term assessments are targeted for additional interventions.

We have taken a different approach to a somewhat similar problem. The problem we are trying to address is how to best assist and support learners during a course when the benefits of big data can be put to work. That is, if we know about the students' academic history, personal history (including demographics), and current activity (such as progress in a course and other related activity pertinent to academic success), what could we do to help? Help would not be for only the struggling student, but for the successful and exceptional students too. The approach we have taken is to construct individualized, personalized advice for students in a large courses on the basis of their academic, personal, and activity profiles (including current progress in the course). We have developed and implemented the Student Advice Recommender Agent (SARA), which generates and delivers an "advice string" to each student each week throughout the term.

Predictive models of student success in the course are computed based on past academic performance and demographic student data. Advice string templates are constructed by instructional experts, focusing on available supports and resources, words of encouragement, and content specific matters. These advice templates are personalized (adjusted/adapted) based upon combinations of student demographic and student activity data. The engineering of advice strings and conditional adaptation is aided by focusing on personas of students who are predicted to fail, pass or excel (as mapped out in [2]). The advice strings are then delivered as learning alerts to each and every student in a course. The advice directs students toward help resources, help or advisory personnel, supplementary course materials, or enrichment activities, as is appropriate.

2 Enhanced Demographics

Beginning in the fall of 2013, the University of Saskatchewan initiated a project to gather enhanced demographic data about incoming freshmen. A 75-item census-style survey is issued annually to students at the time of their first registration to the University. Students are invited to disclose personal information including: goals and aspirations, anticipated and "disappointment threshold" grade point averages; living arrangements; family history of university studies; sources of financial support; expected hours to be spent working for pay, volunteering, studying, engaging in extracurricular activity; disabilities; whether they are supporting dependents; expected major; anticipated advanced or professional degrees sought; level of comfort and connectedness with campus. In addition some short standardized instruments including: a shortened version of Biggs' Study Process Questionnaire [4], the Motivated Strategies for Learning Questionnaire [5] and GRIT-S [3]. Response rates around 60% have been achieved in each of the first two years of this survey.

These data are merged into the University's student data warehouse where information is consolidated from student admissions and recruitment, student grades and academic records, access to academic support services, and learning management system activity.

This rich data repository offers an opportunity to develop comprehensive and relatively accurate predictive models of student academic achievement. The rich data repository also opens possible avenues for inappropriate and prejudicial decisions about students. Data are carefully guarded, identities are encrypted, and personally identifying variables are kept separate from other data. Strict ethical guidelines are followed in making use of the data for student modeling and advising.

3 Temporally Improving Predictive Models

After various data mining attempts using decision trees, regression models, Bayesian networks and naïve Bayes algorithms, predictive models of student academic success in specific courses, overall GPA, retention, and degree completion were derived using student data over the past 5 years. We took a closer look at a face-to-face introductory Biology class that reaches more than a thousand students (mostly freshmen) per year. A sequence of predictive models for final course grades was developed, one model for each week of the Biology course, using the data that would be known as of that week. Models were built with half the students in the 2013 Biology cohort and validated with the other half of the students. We found that log-linear regression models based upon selected demographic data and high school grades could result in good correlations with 2013 course grades (r=0.61). We also found that if LMS activity data and term grades were added in, the grade prediction improves even more as the term progresses. After the course midterm examinations, regression models correlated very highly (r=0.92) with final 2013 course grades.

When these models were applied to the 2014 students, we discovered that the models correlated very well with the 2014 students' grades (r=0.62 at the beginning of term and r=0.82 after the midterm exam). Because of the interventions associated with the introduction of SARA in 2014, we expected correlations between model and final grades might be reduced somewhat – the predictive model applied did not account for changes in 2014. This will be explained in the evaluation section below.

This predictive modeling methodology provides a temporally improving predictive model of student academic achievement in a single course. The model gives a relatively accurate estimate of student success. Factors in the model that offer the greatest degree of predictive power include: actual assessment grades, High school GPA and Biology grades, whether the student was intrinsically or extrinsically motivated, whether the student was a deep versus surface learner, and whether or not the student was the first in their family to attend university.

The methodology also identifies which demographic variables may be considered as risk factors in student retention and success. Using risk factor variables we have constructed a number of personas of canonical successful or less successful students (cf. Brooks & Greer, 2014). Figure 1 shows three sample personas of students.

Student 1	Student 2	Student 3
- rural high school - average admission GPA - first in family at univ - living in univ. residence - high grit score - surface learner tendency predicted to receive D grade	- mature student - has dependents at home - 20 hours per week job - returning - 5 yrs away - deep learner tendency predicted to receive a C grade in course	- exceptional admission GPA - attended top high school - recently settled immigrant - living at home with family - surface learning tendency - aspirations for grad school predicted to receive a B grade in the course

Figure 1: Sample personas of learners

4 CREATION OF ADVICE STRINGS

The personas help our instructional designers, instructors, and academic support specialists construct advice templates that could be tailored for individuals in a particular persona group. Engineering advice strings has turned out to be a fairly difficult activity. Good quality advice for students is highly contextualized – dependent on the time in the term, what is going on in the course, the content being presented, the supports and resources available outside the course, upcoming events and opportunities, news and current events. Good advice also should reflect the student's situation interests and needs, the academic risks they may face, and their determination to survive or excel. For some students useful advice is a message that somebody cares about their success. For others such a message may be a threat and lead to discouragement.

Another challenge with generating advice strings is that predictions may be wrong. Some students predicted to earn an A will not do so. Some students predicted to fail may surprise everyone. We have paid a lot of attention to ensure that our advice to students "does no harm". Advice is framed as a set of positive suggestions, raising opportunities, making reference to resources and supports while offering a supportive and caring tone.

Figure 2 shows the rule for generation of advice in week 3 of our Biology course. The advice strings are written by an subject area learning specialist and they are crafted for stereotypical students who have certain attributes or whose persona has certain features, specified as constraints. The advice string constraints are then interpreted for the attributes of students in the class and a unique advice string is produced for each student. The advice constraints in Figure 2 yield 24 different advice messages. A student will receive the one that best fits their persona. Some weeks there might be only half a dozen different advice messages and some weeks there may be hundreds of different messages. The approximately 1200 students in this Biology course each received a weekly message, tailored as much as possible for their individual context.

We see the messages from SARA as a type of mass personalization. Over the course of a full semester, the collection of weekly advice messages for each of our ~1200 Biology students is nearly unique. That is, the cumulative advice for an individual is likely to be distinct from the cumulative advice given to any other individual in the course. The largest group of individuals who received identical cumulative messages over the term was of size 10. These 10 students were students who were average in every way and for whom we had no enriched demographic data (they did not complete the entry census).

Our students use the BBLearn LMS for access to course materials, lecture recordings, presentation slides, and online quizzes. In order to be sure that students see the weekly advice from SARA (knowing that students tend not to read email), we added an iframe to the course home page where each student sees their weekly "A Note From SARA" immediately as they connect to BBLearn. In addition, an LTI component has been inserted into BBLearn, where all their advice strings for the term so far are available for review, where the advice that SARA gives to others can be browsed (identities hidden), and an opportunity is available to rate the usefulness of or comment upon SARA's most recent piece of advice.

Week: **3**	Order: **0**	Condition: **Predicted GPA < 60%**

Finding balance is important and is often one aspect of academic success that doesn't receive much focus. Regular practice and review of course materials have been proven to help retention of information. This is why completing weekly quizzes, preparing for your labs, and regularly reviewing your textbook and lecture notes will be beneficial to your learning in Biol120! Once you've addressed your academic responsibilities then you can take time for other responsibilities and personal time. Check out this video for additional tips on how to find balance this term.

Week: **3**	Order: **0**	Condition: **Predicted GPA 60% - 80%**

Finding balance is very important and often one aspect of academic success that doesn't receive much focus. Self-awareness, knowing your goals, and self-reflection: these things will be helpful guides in your journey to finding balance. Try to schedule your time so that you are accounting for time to study, time for your other responsibilities, and personal time. Also keep in mind that this schedule will need to be flexible to account for preparation before important deadlines and exams. For additional tips on how to find balance, check out this video.

Week: **3**	Order: **0**	Condition: **Predicted GPA > 80%**

One aspect of academic success that doesn't receive much focus is how to find balance between the time spent on your academics and the time dedicated to your other life responsibilities. Self-awareness, knowing your goals, and self-reflection: these things will be helpful guides in your journey to finding balance. Scheduling time to study is important, but so too is taking time for enjoying other things in life. Check out this video for additional helpful tips on how to find balance this term.

Week: **3**	Order: **1**	Condition: **Low # hours spent studying** (according to survey response and course load)

You indicated on the entry census that you intend to allocate a <i>below average</i> amount of study time on your University courses. Making time to review and study the material is a very important aspect of your overall academic success at University. Setting academic goals and constructing a plan to achieve those goals can be very helpful in guiding how much time you need to prioritize to your studies. Check out these helpful guides to Goal Setting and
 Time Management.

Week: **3**	Order: **1**	Condition: **Medium # hours spent studying** (according to survey response and course load)

You indicated on the entry census that you intend to allocate an average amount of time to studying your University courses. This is a great step towards helping you find balance while also achieving your academic goals! Understanding more about your learning process may help you to use your study time more efficiently - check out this resource on Learning Styles for further information.

Week: **3**	Order: **1**	Condition: **High # hours spent studying** (survey response and course load)

You indicated on the entry census that you would be spending <i>more than</i> an average amount of time studying for your University courses. Did you know that often it's not how long you study, but how efficiently you study, that makes the biggest difference? Using a variety of study methods and taking frequent breaks can help to increase your retention of the material you are studying. To help you use the time you study most efficiently, check out this resource on Learning Styles that might help you understand more about your learning process.

Week: **3**	Order: **2**	Condition: **Working 13-20 hours per week or More and** Predicted GPA < 80%

Working a lot of hours can place extra stress on you, especially as exam times draw near. Be sure you can find a balance between work and your academics. This may mean asking for some time off during heavier study periods. Keep in mind that dropping a few shifts may be more cost and time effective than having to re-take a course!

Week: **3**	Order: **3**	Condition: **Predicted GPA < 60**

Volunteering can provide enriching experiences that can be beneficial to your academic experience and your future career. However, making sure that you are achieving your academic goals needs to be your first priority. After you receive your mid-term grades, if you're on track with success, you may want to consider getting involved with fun and interesting volunteer experiences.

Week: **3**	Order: **3**	Condition: **Predicted GPA 60% - 80%**

One way to achieve balance is to get involved in interesting volunteer opportunities. Volunteering can provide enriching experiences that can be beneficial to your academic experience and your future career. There are lots of ways to get involved on and off campus. To find volunteer positions, check out these volunteering opportunities. If you're interested in gaining valuable experience in a Biology related field, find out about possible volunteering positions by speaking to your Instructor or TA, or contacting the undergraduate's Biology Club.

Week: **3**	Order: **3**	Condition: **Predicted GPA 60% - 80%**

Volunteering can provide enriching experiences that can be beneficial to your academic experience and your future career. Find volunteer opportunities that might be of interest to you by checking out these volunteering opportunities. In addition, if you're interested in gaining valuable experience in a Biology related field, find out about possible volunteering positions by speaking to your Instructor or TA, or contacting the undergraduate's Biology Club.

Figure 2. Different Advice Strings generated by SARA in week 3

5 Evaluation of SARA

As with many educational interventions initiated in a large course, controlled experimentation is difficult to implement. In the freshman Biology course of 2014 many factors remained the same as in prior offerings. The instructional team, the course objectives and evaluation rubric, the laboratories, and the types of students remained much the same. Key differences in 2014 included a new requirement of weekly online quizzes, the introduction of SARA's advice, and the addition of optional-attendance peer-led study groups.

A remarkable difference in achievement was detected when comparing unadjusted final grades in 2014 against the two previous years. The mean course grade increased by 2.57 percentage points (t-test $p<.00001$). More remarkably, the number of D and F grades in 2014 decreased by 25.3% over the previous year and the number of A grades increased by 28.3% over the previous year. In the ~1200 student course, 100 fewer students scored D and F grades in 2014 than in previous years, indicating a potential boost in student retention.

Students' predicted grades for 2014 were also compared against their observed grades. The predicted model at the beginning of term expected a class average 3.8 percentage points lower than the observed average. After the midterm, the improved predictor expected a class average that turned out to be 1.8 percentage points lower than what was observed. These significantly lower predictions may indicate that the predictive model was lacking. Of course we hoped that the model would underpredict grades if there was indeed a positive effect due to the new teaching interventions.

The fact that correlations between predicted and observed grades remained comparable to the prior year indicates that there was a general lift to grades in 2014 (no significant change in slope of the correlation line). There was clearly an across the board increase in grades in 2014. Instructors believed this was due primarily to the introduction of the weekly quizzes.

An end of term questionnaire was given to determine students' reactions to changes made in the class. Forty four percent of the students completed the survey but nearly half of those students chose to remain anonymous so their responses could not be linked with grades and demographics. Students who completed the survey and gave their identities tended to have somewhat higher grades than those who did not give identities but the patterns of responses among those who did and did not reveal identities were not significantly different. Among the survey items students were asked whether they read the weekly advice from SARA or ignored it and whether or not they appreciated SARA's advice. Only 1/3 of students said they paid attention to SARA's advice, while 40% of the students said they appreciated receiving advice from SARA. Neither students who paid attention to SARA's advice or students who appreciated receiving advice from SARA showed differences in unadjusted final grades or in differences between expected (pre-midterm) and actual grades as compared to students who ignored SARA. Similarly there was no difference in grades between those who chose to participate in the peer-led study groups and those who did not.

Perhaps in an ideal world every student should embrace and act upon SARA's advice. But it is important to realize that even if only a few students listen to and are helped by SARA's advice, and that the advice is helpful, a significant shift in achievement can (and did) occur.

One important measure related to SARA's advice did show promise. As stated above, the post-midterm predictor, correlated with unadjusted final grades at 0.82. We examined more closely the error in this predictor (the difference between actual unadjusted final grade and the post-midterm prediction). Students who regularly read the weekly advice from SARA scored significantly higher than the predictor (4.6 percentage points), while students who did not read SARA's advice scored very near the predicted grade. This could be taken to mean that student's who regularly read SARA's advice had achievement levels higher than was expected.

Given our study and its limitations, we cannot directly attribute the overall achievement improvements in this offering of the course to SARA alone. Improvements could have been due to the weekly quizzes, which every student was required to complete, or the study groups, or SARA or some combination. The Biology instructors were very happy with the outcomes in their course. They were convinced that more learning occurred than in past years and that the combination of interventions was a great success. Yet to more fully understand the impact of SARA, we are contemplating dropping the SARA advice from the Biology course next year and watching to see if the increase in level of achievement persists when only mandatory quizzes and optional study groups are in play. In order to continue our research into SARA we hope to expand SARA's reach into other large freshman STEM courses in the upcoming year (perhaps Physics, Chemistry or Engineering).

6 Conclusions

SARA provides a scalable advice personalization environment in large university courses. In our first offering of a large course using SARA, student achievement improved over previous years and students on average achieved significantly higher grades than our predictive model (based on prior years' features) would have expected. There is some evidence that the improvement could be caused by SARA's weekly advice, but further research is needed to confirm such a claim.

One of the persistent dangers that comes along with predictive modeling of learners is the possibility of prejudicial treatment that may bring negative consequences to some learners. When instructors come to know, for example, that a particular student has a very low probability of successfully completing a course, the instructor may decide to minimize help and support for that learner, turning more attention toward those with higher likelihood of success. Likewise an instructor may (consciously or subconsciously) privilege the students who would most likely pursue advanced courses or graduate programs. Of even greater concern, predictive models may be used to shape enrolment, streaming, and admission policies. It is important for instructors and especially admissions officers to understand that probabilities associated with predictive estimates make it highly inappropriate to assume the predictive tendencies of a group will apply to any particular individual in the group.

Finally, it is important to remember that the goal of improved learning is to make gains in academic achievement. If using predictive models of achievement based on past cohorts to inform decisions about future cohorts, one must be prepared to accept that the predictive models may be a little less accurate than one might like. This paradox associated with modeling the state of learners as they consciously and steadily try to move beyond their current learning state is not new. The temporal dimension of predictive models that takes into account innovations or interventions in teaching and academic support is vital to our growing understanding of the learning process for distinct individual learners.

Acknowledgements

We would like to recognize the contributions, creativity, and willingness to flex of the three faculty members who taught this Biology course in 2014: Susan Kaminskyj, Ken Wilson and Jorge Chedrese.

References

1. Arnold, K. E., & Pistilli, M. D. (2012, April). Course Signals at Purdue: Using learning analytics to increase student success. In *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge* (pp. 267-270). ACM.
2. Brooks, C., & Greer, J. (2014, March). Explaining predictive models to learning specialists using personas. In *Proceedins of the Fourth International Conference on Learning Analytics And Knowledge* (pp. 26-30). ACM.
3. Duckworth, A. L., & Quinn, P. D. (2009). Development and validation of the Short Grit Scale (GRIT–S). *Journal of personality assessment*, *91*(2), 166-174.
4. Fox, R. A., McManus, I. C., & Winder, B. C. (2001). The shortened Study Process Questionnaire: An investigation of its structure and longitudinal stability using confirmatory factor analysis. *British Journal of Educational Psychology*, *71*(4), 511-530.
5. Pintrich, P. R., Smith, D. A., García, T., & McKeachie, W. J. (1993). Reliability and predictive validity of the Motivated Strategies for Learning Questionnaire (MSLQ). *Educational and psychological measurement*, *53*(3), 801-813.