Open Academic Analytics Initiative (OAAI)
Initial Research Findings

Dr. Eitel Lauria, Dr. Erik Moody, Mr. Sandeep Jayaprakash, Ms. Nagamani Jonnalagadda, Mr. Josh Baron

DISCLAIMER: This report on our initial research findings is intended for informational purposes and thus details related to methodology, limitations, and implementation challenges have not been included. These details will be provided in future full length reports as well as peer-reviewed articles. Questions or requests for these details can be sent to Josh Baron, OAAI Principal Investigator, at the following email: Josh.Baron@marist.edu.

Marist College

3399 North Road

Poughkeepsie, New York 12601

9/7/2012

Executive Summary

During the spring 2012 semester, the Open Academic Analytics Initiative (OAAI), led by Marist College, successfully deployed an open-source Learning Analytics solution, developed by the project during the fall, at two community colleges (Cerritos College and College of the Redwoods) and one Historically Black College and University (HBCU) (Savannah State University) as a means to further research in this emerging field. Our spring pilots involved a total of 1379 students, 67% of whom were considered low-income students, who were enrolled in introductory-level courses with, generally, three sections each being taught by the same instructor (e.g. BIOL 101 Section 1, 2 and 3). Each course section was then assigned to either a control or one of two treatment groups, thereby standardizing the instructional delivery to the extent possible across all three. Students in the two treatment groups who had been identified by our predictive model, which uses student demographic, aptitude and course management system usage data, as being likely to not complete the course received interventions designed to help them succeed.

This report discusses our initial results from our spring pilots with regards to our two primary research areas: the "portability" of predictive models from one academic context to another and the effectiveness of different intervention strategies. In summary, we have found the following to date:

- The predictive model developed using student data from Marist College was very similar with regards to the predictive elements (e.g. cumulative GPA) and correlation strengths as the predictive model developed at Purdue University by Dr. John Campbell.
- The accuracy of the predictive model built using Marist student data performed considerably higher than sheer randomness when deployed at both community colleges and HBCUs which exceed our expectations.
- We have found a **_statistically significant difference_** between mean course grades when comparing all students in our two treatment groups (awareness and OASE) to our controls.
- In addition, we have found a **_statistically significant difference_** with regards to "content mastery" (a C grade or above) as well as withdrawal rates between our combined treatment groups and our controls.

These findings, which are detailed on the following pages, are noteworthy as they indicate that predictive models used in Learning Analytics are more "portable" then initially anticipated as well as the fact that our intervention strategies have been effective in improving student outcomes.

Overview of Initial Portability Research Findings

As discussed in our original grant proposal, many questions exist in the emerging field of Learning Analytics with regards to the degree to which predictive models which are built based on data from one particular institution type and student population can be effectively deployed in academic contexts where the institution type and/or student population differ significantly. The degree to which predictive models may be "portable" has major implications for the scaling of learning analytics across multiple institutions and even higher education itself. The OAAI has researched this issue of portability in two phases, first looking at the degree to which a model built using data from Marist College compared to the original model built at Purdue University (which lead to the development of Course Signals) and then, second, deploying the Marist-developed model at community colleges and HBCUs and investigating the accuracy of the model in these different academic contexts. Results from these two research efforts, which are detailed below, have shown that predictive models appear to be more portable than initially anticipated.

Phase One Portability Research Results: Comparing Marist Predictive Model to Purdue Model

During the summer of 2011 the analytics research team at Marist College developed a predictive model using the same student demographic (e.g. age), aptitude (e.g. SAT scores), and course management system usage (e.g. number of course visits) that were used by Dr. John Campbell at Purdue University in his original dissertation research except our data was from Marist students and based on usage of the Sakai[1] Course Management System (CMS). Although Marist College and Purdue University differ in obvious ways (e.g., institutional type, size, and instructional approaches) they do share a number of similarities which are particularly pertinent to this study. These include percentage of students receiving federal Pell Grants (Marist 11%, Purdue 14%), percentage Asian/Black/African American/Hispanic students (Marist

Table 1. Correlations between course grades and CMS data

Undergraduate CMS event frequencies		Course Grade	
		Marist Fall 2010 N=18968	**Campbell (2007)** N=27276
Sessions Opened	Correlation	0.147	(no values reported)
	Significance	0.000(**)	
	N	11195	
Content Viewed	Correlation	0.098	0.112
	Significance	0.000(**)	0.000(**)
	N	7651	19205
Discussions Read	Correlation	0.133	0.068
	Significance	0.000(**)	0.000(**)
	N	1552	7667
Discussions Posted	Correlation	0.233	0.061
	Significance	0.000(**)	0.000(**)
	N	1507	7292
Assign. Submitted	Correlation	0.146	0.163
	Significance	0.000(**)	0.000(**)
	N	3245	4309
Asssmts Submitted	Correlation	0.161	0.238
	Significance	0.000(**)	0.000(**)
	N	1423	4085

(**) Significant at the 0.01 level (2-tailed)
Marist data uses ratios over course mean instead of frequencies

---

[1] Purdue was using WebCT/BlackBoard in their research.

11%, Purdue 11%), and ACT composite 25th/75th percentile (Marist 23/27, Purdue 23/29).

Once the Marist predictive model was developed, we compared the elements of the model that were statistically significant with regards to correlation with student grades (as was done at Purdue) as means to understand to the degree to which the models differed. In general, we found the same statistically significant elements as Purdue with similar correlation strengths (see Table 1 for example of CMS events). The OAAI analytics research team published a paper (*Mining Academic Data to Improve College Student Retention: An Open Source Perspective*[2]), related to this initial research on "portability" which was presented at the 2012 international Learning Analytics and Knowledge (LAK) conference.

Phase Two Portability Research: Deployment of Marist Model in Different Academic Contexts

After the phase one research efforts, additional time was spent improving the initial predictive model through additional machine learning techniques and the introduction of additional data elements (e.g. gradebook data from the CMS). The enhanced predictive model was then evaluated using a Marist test data set which had been excluded during the development of the model. Table 2 provides a summary of the outcomes from this evaluation (a total of ten trials were run but the table only includes the means from these tests).

This model was then deployed as part of our course pilots during the spring semester with two community colleges and a HBCU which represent vastly different educational contexts as compared to Marist. For example, Savannah State has a 94%

Table 2: Prediction Analysis using with Marist Test Data

Classification Algorithm	Performance Metrics	Mean	SE
SVN / SMO	Accuracy	85.96%	0.86%
	Recall	81.24%	4.55%
	Specificity	86.31%	1.24%
	Precision	31.81%	2.16%
C5.0 / J48	Accuracy	86.82%	1.14%
	Recall	65.45%	3.85%
	Specificity	88.31%	1.48%
	Precision	31.34%	2.52%
Logistic Reg.	Accuracy	86.07%	0.92%
	Recall	80.26%	5.18%
	Specificity	86.53%	1.35%
	Precision	33.01%	3.44%

Black non-Hispanic student population with 67% of receiving Pell Grants, College of the Redwoods' student population has 22% minorities and 60% are receiving Pell Grants and Cerritos College's student body is 41% Hispanic and 45% of the students are receiving Pell Grants.

---

[2] Eitel J. M. Lauría, Joshua D. Baron, Mallika Devireddy, Venniraiselvi Sundararaju, and Sandeep M. Jayaprakash. 2012. Mining academic data to improve college student retention: an open source perspective. In *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge* (LAK '12). ACM, New York, NY, USA, 139-142. DOI=10.1145/2330601.2330637 http://doi.acm.org/10.1145/2330601.2330637

At the conclusion of the spring semester, a similar evaluation was completed to determine how well the model performed when deployed in these different academic contexts. This evaluation included assessing the model's performance at three points during the semester (25%, 50% and 75% of the semester completed), which correspond to when Academic Alert Reports were provided to instructors, to evaluate how the model's performance improved as more CMS and gradebook data was available. Table 3 provides a summary of the results of this evaluation. Looking at the accuracy of the model (percentage of students that were correctly identified), just as one indicator, it is clear that the results from the predictive analysis were considerably higher than sheer

Table 3: Prediction Analysis from Spring Pilots						
College	% of Semester Completed	# of students	Accuracy	Recall	Specificity	Precision
Savannah	25%	504	66.96%	70.76%	64.64%	55.00%
	50%	504	71.52%	78.22%	67.56%	59.41%
	75%	504	77.75%	72.53%	80.94%	69.84%
Cerritos	25%	502	59.13%	69.23%	56.31%	30.73%
	50%	601	70.92%	66.14%	72.51%	44.44%
	75%	649	74.77%	74.42%	74.88%	47.76%
Redwoods	25%	195	70.50%	86.27%	61.36%	56.41%
	50%	195	79.86%	72.55%	84.09%	72.55%
	75%	195	79.14%	72.55%	82.95%	71.15%

randomness. For example, when one compares the accuracy at the "75% of the semester completed" point, which ranged from 75% to 79%, to the accuracy of the model when tested with Marist data (which was historical and thus represented 100% of the semester completed), which was 86-87%, we find only a 6-10% difference. Given that we expected a much larger difference between how the model performed when tested with Marist data and when deployed at community colleges and HBCUs, this was an encouraging finding.

Portability Research Discussion

Our findings from Phase One and Two of our portability research seem to indicate that predictive models that are developed based on data from one institution may be scalable to other institutions, even those that are different with regards to institutional type, student population and instructional practices. We believe this unexpected finding may be the result of the specific elements of the predictive models which have shown to be the most powerful predictors. The elements that are most predictive of student outcomes, which were also identified by Purdue in their research, are cumulative GPA and data from the CMS' gradebook. Given that these two elements are such fundamental aspects of academic success it is not all that surprising that the predictive model has fared so well across these different institutions.

If this explanation is correct, it does point to the importance of instructors using the gradebook within their CMS if they wish to take advantage of learning analytics. It also indicates that models may not "port" well to institutions where cumulative GPA is not available (e.g. non-credit training programs) or if the student population is entering an institution and thus GPA is not yet available. Thus, although our initial findings are encouraging with regards to portability, important questions remain with regards scaling up models across more diverse academic settings.

Overview of Initial Intervention Research

OAAI has also conducted research on two different intervention strategies which were deployed to "at risk" students as means to help them succeed in the course. One of these strategies is "awareness messaging", which is closely aligned with the approach taken by Purdue's Course Signals project, which entails the instructor sending a message to the "at risk" student noting their concern over the student's academic performance and then suggesting specific steps the student should take to improve (e.g. meet with a tutor, attend a study group session, etc.). The text of these messages, outside of what the instructor suggests as means to improve, was standardized across all of the treatment groups. The other intervention strategy is referred to as the "Online Academic Support Environment" or OASE and it parallels the "awareness messaging" strategy in that students receive the same initial message noting concern but are then invited to join an Sakai-based online support site in which they are given access to Open Educational Resources (OER) instructional materials (e.g. Khan Academy videos, Flat World Knowledge textbooks, etc.). In addition to these materials, they are provided with a range of mentoring from peers and professional support staff. As with the "awareness messaging", the text of the messages is standardized across instructors and the text becomes increasingly serious in tone as students receive their second and third message.

Intervention Research: Impact on General Student Academic Success

 We began our analysis by comparing, through a one-way ANOVA, the overall academic success, as measured by average course grades, between the two treatment groups and the controls.   The results indicate a statistically significant difference (see Table 4) which we feel is impressive as the groups include all students in the classes assigned to the three groups regardless of whether they were identified by the model as being "at risk" or not.

Table 4: Average Course Grade Analysis

	Awareness	OASE	Control
Mean	77.47%	77.50%	75.14%
n	369	348	285

$F_{(2,999)} = 3.065$, $p = .047*$

We then refined this analysis by comparing, also through a one-way AVNOA, what we believe represents those students who were the most "at risk". In the treatment groups ("awareness messaging" and OASE), we consider only students who had received at least one intervention based on any of the three Academic Alert Reports that were posted during the semester. For controls, which by definition did not received interventions, we selected those students who had been identified as having an average "risk level" of

Table 5: Refined Course Grade Analysis

	Awareness	OASE	Control
Mean	72.05%	72.43%	65.38%
n	193	179	79

$F_{(2,448)} = 8.484, p = .000*$

three or higher across all three Academic Alert Reports. Students were categorized, based on scoring by the predictive model, as having: "no risk" (1), low risk (2), medium risk (3) or high risk (4) with regards to their likelihood of not completing the course successfully[3] . In this analysis (see Table 5) we found even greater statistical significance which may indicate that the interventions were of particular benefit to those "at risk" to not succeed academically.

Table 6: Low Income Student Analysis

	Awareness	OASE	Control
Mean	76.85%	76.61%	74.14%
n	254	242	198

$F_{(2,691)} = 2.601, p = .075$ N.S

 Finally, we also analyzed the impact the interventions had on "low income[4]" students with regards to their average course grades. In this case we did not find a statistically significant difference although the trend seems to indicate that "low income" students in the treatment groups performed better than those in the control groups. Since the number is close to being significant (see Table 6) we are particularly interested in performing analysis on a larger sample containing Pell Grant students which we hope to collect during the fall pilots.

Intervention Research: Impact on Content Mastery

We also examined, through a Chi-square analysis, the impact our interventions had on student "content mastery" (i.e. received a C or better in the course). We performed this analysis by comparing the controls to

Table 7: Students Who Mastered Content (C or above)

		Group	
		Control	Treatment
Content	Yes	179	308
Mastery	No	50	44

Chi-Square (1) = 8.913 p= .003*

---

[3] More specific probabilities associated with these risk levels can be provided upon request.
[4] Low-income students were defined as students who were receiving Pell Grants.

Grading Scale	
Letter Grade	Percentage Grade
F	0 - 55%
D	56% - 65%
C	66% - 75%
B	76% - 85%
A	86% - 100%

both treatment groups combined[5] to increase our sample size. Again, we found a statistically significant difference (see Table 7) indicating that our treatment groups achieved better learning outcomes than those in our control groups.

Intervention Research: Impact on Withdrawal Rates

Finally, we examined the impact our interventions had on withdrawals by comparing withdrawal rates among those students identified as being "at risk" on the first Academic Alert Report in our control groups to the withdrawal rate of the combined treatment group (both "awareness messaging" and OASE). We chose to focus on those students identified during the first Academic Alert Report as we have additional evidence that interventions early in the semester have the largest impact on student outcomes and come at a time when withdrawal is possible without incurring the maximum penalties (e.g. no tuition refund).

Table 8: Analysis of Withdrawal Rates

		Group	
		Control	Treatment
Withdrawal	Yes	29	78
	No	290	425

Chi-Square (1) = 7.097 p= .008*

Based on the Chi-square analysis (see Table 8) we have identified a statistically significant difference between the two aforementioned groups which indicates that students in our combined treatment group were *more likely* to withdraw than those in our control groups. Although this finding may at first seem to be a negative outcome, it is consistent with findings at Purdue and may be an indication that students who receive interventions are withdrawing earlier in the course (as opposed to remaining enrolled and failing) than those who do not. While we would much prefer that students complete their course successfully, withdrawing, particularly before incurring major penalties, is preferable over failing and thus we see this as a potentially positive result. We are working to explore at what point in the semester students did withdraw as means to investigate this further.

Future Research Agenda

Based on our initial findings, we have begun to identify additional research interests for the remaining duration of the grant (which will conclude in January 2013) as well as what we believe is an important national research agenda items related to Learning Analytics. Both of these are briefly discussed below.

---

[5] We believe this is valid as both treatment groups received the same message noting their instructor's concern.

Fall 2012 Research Interests

We will be conducting additional pilots again during the fall 2012 semesters at the three original institutions as well as a new institutional partner, North Caroline A&T (an HBCU). These pilots will provide us with an additional dataset that should contain several hundred more students over what we had from the spring pilots. In addition to continuing to refine our analysis of the spring dataset, we hope to use the fall data to:

- Confirm our findings from the spring data set.
- Increase our overall sample sizes (n) which we believe will allow us to identify additional correlations which are statistically significant.
- Look at our impact on persistence rates which require two semesters worth of data.
- Explore how differences between course subject matter or size might affect the impact of our interventions and/or be a factor in our predictions.
- Look for relationships between student responses' to questions provided by the National Survey of Student Engagement (NSSE) (which were incorporated into our OAAI Student Survey) and the receipt of an intervention. We would also like to look at how NSSE responses relate to student outcomes to explore the possibility of using student responses to NSSE questions in our predictive model.
- We are also planning on conducting interviews with select instructors and students to identify strengths and weakness of the interventions. These insights will help us identify ways to refine our intervention approaches.

Potential Items for a National Learning Analytics Research Agenda

Based on our work to date, our research team has begun to discuss and identify areas of research that we believe will be important to the field of Learning Analytics as it begins to be deployed more widely. These research questions, which are outlined below, could form the basis for a national research agenda in this new and emerging field.

1. *Does Learning Analytics allow us to identify students "at risk" to not complete a course that the typical instructor would miss?*

    Learning Analytics solution are often seen as ways to improve student success in large lecture-style courses (100+ students) in which it can be very difficult for an instructor to identify, early on in the course, which students may not succeed. As Learning Analytics is deployed at

institutions with smaller classes sizes, as was the case with many of the OAAI course pilots, it will be important to understand the "value add" that Learning Analytics brings over what an instructor is capable of doing on their own. Although many of the instructors in our pilots have noted with us that they have found the identification of "at risk" students very helpful, it remains unclear to us if they would have identified the same students as our model if they attempted to do so on their own. Thus, we believe it will be important to conduct research studies in which "instructor predictions" are compared to "model predictions" as means to better understand the answer to this question.

2. ***What are the characteristics of students who seem to have "immunity" to the treatment (those who got interventions but never improved) and those who were effectively treated after just one intervention?***

   From our initial research we have found that student seem to fall into one of two broad categories, those who improved after receiving just one "treatment" or intervention and those who did not improve regardless of the number of treatments received. Very few students who did not improve after the first intervention went on to improve after the second or third. Our theory is that some students respond very well to the "treatment" and thus improve after just one intervention while other seem "immune" to the "treatment" and do not improve regardless of how many treatment their receive. Understanding why this is the case and what characteristics are associated with these two categories of students would help us better understand how to most effectively deploy interventions.

3. ***How portable are predictive models that are designed for one type of course delivery (e.g. face-to-face) when they are deployed in another delivery format (e.g. fully online)?***

   We are particularly interested in exploring the issue of portability with regards to face-to-face and fully online programs given how much more CMS usage takes place in the later mode of instruction. It may be that models that are developed based on face-to-face courses do not "port" well to fully online or at least that such models could be significantly improved if they were customized for fully online programs.