Seven-year longitudinal implications of wheel spinning and productive persistence

Seth A. Adjei[1], Ryan S. Baker[2], Vedant Bahel[3]

[1]Northern Kentucky University, Highland Heights, KY 41099, USA
[2] University of Pennsylvania, Philadelphia, PA 19104, USA
[3] G H Raisoni College of Engineering, Nagpur, India
adjeis1@nku.edu

Abstract. Research in learning analytics and educational data mining has sometimes failed to distinguish between wheel-spinning and more productive forms of persistence, when students are working in online learning system. This work has, in cases, treated any student who completes more than ten items on a topic without mastering it as being in need of intervention. By contrast, the broader fields of education and human development have recognized the value of grit and persistence for long-term outcomes. In this paper, we compare the longitudinal impact of wheel-spinning and productive persistence (completing many items but eventually mastering the topic) in online learning, utilizing a publicly available data set. We connect behavior during learning in middle school mathematics to a student's eventual enrollment (or failure to enroll) in college. We find that productive persistence during middle school mathematics is associated with a higher probability of college enrollment, and that wheel-spinning during middle school mathematics is not statistically significantly associated with college enrollment in either direction. The findings around productive persistence remain statistically significant even when controlling for affect and disengaged behavior.

Keywords: Wheel-spinning, Grit, Productive Persistence, College Enrollment.

1 Introduction

Grit, the combination of persistence and passion, is important to both learning and life outcomes [11, 12]. The benefits of grit can span across several years [12]. Recent work has argued that the persistence component of grit is more important to student outcomes than passion [9]. Ultimately, the ability to work hard towards a goal – and not give up even in the face of serious challenge – appears to be a key part of life success.

At the level of courses and programs, the learning analytics literature recognizes the value of persistence. There has been considerable research on studying stopout/ dropout, quitting a course or program prior to completion. There has also been a proliferation of research on models that can detect that a learner is at risk of stopping or dropping out [10, 14, 19], as well as work on understanding the factors that lead students to stop out or drop out [29, 31, 35]. This work underpins the development of systems used to prevent stopout and dropout, which in many cases have led to better outcomes for learners, where they persevere to successful completion [2, 21, 33].

Curiously, however, the more micro aspects of the learning analytics literature – looking at student behavior within a problem set, for instance -- have largely treated perseverance as a problem rather than as a strength. The term “wheel-spinning” has been applied to this behavior, viewing persistence as being the same thing as making no progress [4] The initial definition of wheel-spinning proposed by Beck and Gong [4] treated any student who completed ten mathematics problems on a single skill without mastering the skill as wheel-spinning. This has continued as the most common definition of wheel-spinning (though some work has hand-labeled wheel-spinning – i.e. [20]); the preponderance of research on wheel-spinning has not differentiated between students who persevere but eventually succeed, and students who persevere and never succeed (i.e.[7, 13, 15]). However, Kai and colleagues [17] have noted that many students continue to make progress, and obtain mastery, even after completing ten problems. Käser and her colleagues [18] argue that a student should only be considered wheel-spinning if they are no longer making progress in terms of a knowledge model.

Following on work by Kai et al. [17], we adopt a different view on persistence during the learning process, separating persistence that ultimately leads to success (termed “productive persistence”) from persistence that never does (which we term “wheel-spinning”, somewhat in contrast to the use in [4]).

2 Research Questions

In this paper, we investigate the relationship between measures of learner persistence in the ASSISTments system, used in middle school, and a longitudinal outcome, college enrollment. We measure three indicators of persistence (or the lack thereof):

- Wheel-spinning: The student persists but never succeeds in mastering the skill
- Productive persistence: The student persists and eventually masters the skill
- Quitting: The student does not persist and quits the skill without mastery

We hypothesize that productive persistence will be associated with positive longitudinal outcomes, whereas wheel-spinning and quitting will be associated with negative longitudinal outcomes, albeit for different reasons.

3 The ASSISTments System

In this study, we examine the effects of learners’ persistence on college enrollment within the ASSISTments learning system. ASSISTments [16] is a free online learning system that several thousand middle school teachers in the United States use to assign math homework to their students. On average, over 35,000 students across the United States use this system on a daily basis, solving approximately 323,000 questions a day.

The system provides sets of prepackaged questions or problems, called problem sets, grouped by math topics/skills. Each problem set is made up of a set of problems that are tagged with at least a single math skill, where the skills are sourced from the United

States Common Core State Standards for Mathematics. [8] Many of these problem sets, referred to as “skill builders”, are mastery based, where the system continues giving the student problems until they demonstrate mastery. The mastery criterion (i.e. specific number of questions correct in a row required for mastery) is predetermined for each problem set, although it can be adjusted by the teacher. In general, the mastery criterion is 3 problems correct in a row. If a student completes ten problems on a skill without reaching mastery, the system asks them to take a break on the skill until the next day.

Many ASSISTments problems, including all skill builder problems, have features that allow students to seek help if they experience difficulties as they answer the questions. Figure 1 shows a sample problem presented in the ASSISTments tutor. The tutor gives students access to help-seeking features, including hints and scaffolding questions. Hints are short simple statements or clues about the question that help guide the student through the solution to the problem, for instance by explaining the knowledge component(s) required to solve the problem or providing the formula required to solve the problem. Hints become available in increasing order of specificity, with the final hints usually providing the answer to the question (referred to as the bottom-out hint.) Scaffolding questions [26] are created based on the original math problem. The original problem is broken down into smaller, less difficult steps, designed to be answered in a linear sequence. The answer to the final scaffold question is the answer to the main question. Within skill builders, for each main question, students are provided with both of these help features and can choose either. If a student seeks help on a question, the system can either (according to the teacher’s preference) mark the question incorrect or give the student partial credit based on the number of help-seeking steps they sought [23, 30, 32]. Students’ work within ASSISTments is logged, including the problem identifier, the correctness of the student’s response, the type of help the student sought, the number of hints the student requested, as well as the time (in seconds) the student spent in answering the question.

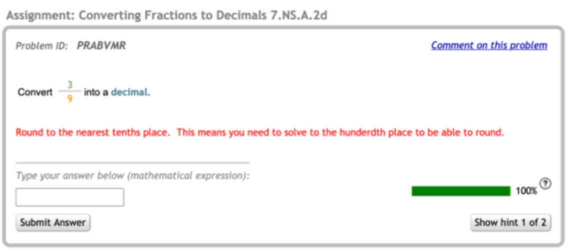

Fig. 1. A sample ASSISTments problem displayed in the tutor. The system presents the student an opportunity to ask for hints and to attempt the question multiple times.

4 Data

4.1 ASSISTments Log

In this paper, we analyze a publicly available ASSISTments dataset, the dataset used in the ASSISTments Longitudinal Data Challenge [25], selected because of the availability of longitudinal outcome data. The overall data set consisted of data from 1,709 students – we analyze the subset of students who completed skill-builder assignments. This smaller dataset was comprised of the problem response logs of 236 students who completed a total of 431 skill-builder assignments in the ASSISTments system. The logs were collected between the 2004 and 2007 school years. On average each student either started (and/or completed) 18 skill-builder assignments, generating a total of 25,159 logs across a total of 5,979 problems.

4.2 College enrollment data

College enrollment records were collected for the 236 students whose logs we examine in this paper (and for the rest of the students in the data set as well). Each student's enrollment record includes a binary feature that indicates whether they enrolled in college. This data was collected through the National Student Clearinghouse, a database on college enrollment, and made available to researchers through the ASSISTments Longitudinal Data Challenge. A full discussion of the collection of this data is given in [28]. Data was collected for college major and post-college job as well [25], but will not be analyzed in this paper due to the smaller sample size.

4.3 Affect and Behavior Data

The data set included estimates of students' affect and behavior as they answer the complete the mastery-based skill builders in ASSISTments. The affective states for which the data set had estimates are boredom, engaged concentration, confusion and frustration. The behavioral estimates are of two forms of students' disengagement as they complete the assignment: gaming the system and off-task behavior. These estimates were originally generated (see [24]) using a two-step process. In the first step, observers trained in the BROMP protocol and HART android app [22] recorded observations of student affect and disengaged behaviors for a small sample of students. In the second step of the process, the observations were used to create models inferring affect and disengaged behaviors from only problem logs. These models were then used to calculate estimates of students' affective and behavioral states for the unseen ASSISTments problem logs. Each affect and behavior estimate was scaled from 0 (0% probability) to 1 (100% probability) by the models.

5 Feature Generation

In order to study how persistence was associated with college enrollment, we created features from the ASSISTments log data to represent different aspects of persistence. The following sections describe each of the features in detail.

5.1 Mastery Speed

In the context of ASSISTments, Xiong et al. [34] defined mastery speed as the number of questions a student answers prior to achieving mastery of a skill. For instance, a student who obtains three right in a row on the 7th, 8th, and 9th problems in a mastery-based assignment will have a mastery speed of 9. (The student will not have achieved the mastery criterion prior to the 9th question) We calculated this feature for each student/assignment pair. A student who already knows a skill will achieve mastery (according to the system) by obtaining correct answers on the first 3 problems in a row. However, students can also be deemed to have mastered the skill if they answer the very first question in the assignment correctly without any help (i.e. hints or scaffolds). These students are deemed to have "tested out" of the assignment. In such cases, the mastery speed for such a student-assignment pair was counted as 1.

5.2 Wheel-Spinning

In this paper, we adopt a definition of wheel-spinning similar to the definition used by Kai and colleagues [17]. As ten problems completed (for persistent) and three-in-a-row correct (for mastery) match the operationalizations used in the ASSISTments system itself, we adopt these definitions. We define a student as wheel-spinning if they complete ten problems without reaching mastery and never reach mastery on that skill. If a student completes ten problems without reaching mastery but then masters the skill, we define them instead as a separate category, persistent-mastered. To reiterate, a student is deemed to be persistent in a skill builder assignment if they are unable to achieve mastery by the 10th question in the assignment – if their problem count for the given assignment is at least 10. The ASSISTments system generally stops the student from being presented with additional questions in an assignment on the same day if they have not mastered the skill by the 10th question (unless the 10th question is answered correctly after the 9th question was also answered correct). Whether or not they eventually master the skill determines whether they are treated as persistent-mastered or wheel-spinning.

5.3 Persistence-related Features

We categorize students in terms of their mastery and persistence in terms of four behaviors, each of which is expressed in a data feature as the percentage of skills where the students demonstrated that behavior: *Persistent-Mastered*, *Wheel-Spinning*, *Quickly-Mastered*, and *Quit*, shown in Table 1. These features are also described below:

(i) **Quickly-Mastered**: The percentage of assignments in which the student mastered the skill in ten or fewer problems.

(ii) **Persistent-Mastered**: The percentage of assignments in which the student attained mastery of the skill builder after more than ten problems (referred to as "productive persistence" in Kai et al. [17]).

(iii) **Wheel-Spinning**: The percentage of assignments in which the student completed more than ten problems and never attained mastery.

(iv) **Quit**: The percentage of assignments in which the student neither attained mastery nor was persistent – they quit working on the problem set prior to the tenth problem. Also referred to as "early stopout" [6].

Table 1. The four categories of persistence and success analyzed in this study.

	Persistent (> 10 problems)	**Non-Persistent (<=10 problems)**
Mastered	Persistent-Mastered	Quickly-Mastered
Did not master	Wheel-Spinning	Quit

5.4 Other Mastery-related and Performance features

In addition to the above-mentioned features, we examined the following high-level features to understand their interaction with college enrollment and better clarify our findings: the percentage of assignments in which the student was persistent (the student exceeded ten problems, irrespective of whether they achieved mastery or not), the percentage of assignments that the student mastered, and the average percent correct across all questions they were presented across all assignments.

5.5 Aggregation of features

As stated earlier, on average each student completed 18 mastery-based assignments. In order to be able to compare the features across these assignments to college enrollment, we aggregated all the features described above such that each student record will have a single average value for each of the affect, behavior and ASSISTments-related features. For example, instead of using a single mastery speed for an assignment in our models, we used the average mastery speed across all assignments.

6 Analysis Plan

To research the effects of wheel-spinning on college enrollment, we conducted three analyses, much as in [28]. In the first analysis, we looked at the relationship between each of the features discussed above, taken individually, and college enrollment. For each feature, we conducted a two-sample two-tailed t-test, assuming equal variances, comparing the value of each feature between students who attended college and students who did not attend college – for example, was wheel-spinning more frequent among students who eventually attended college, or students who did not eventually attend college? Cohen's D effect sizes were used to assess the magnitude of the differences. Given that we ran 18 tests, a Benjamini & Hochberg [5] post-hoc control was used. Benjamini & Hochberg's procedure, a false discovery rate procedure, attempts to ensure that the overall rate of Type II false discoveries (non-effects treated as effects; false positives) remains at the 5% level. In doing so, Benjamini & Hochberg's procedure avoids the over-conservatism and inflated Type II error rate that the Bonferroni correction is known for, while avoiding the inflated Type I error rate that occurs when multiple tests are run and no post-hoc adjustment is used. Within the Benjamini & Hochberg correction, different tests are assigned different alpha values (which the p value must be below to reach statistical significance), based on both the overall number of tests run and the number of tests that have lower p values than the current test.

In the second analysis, we created a single logistic regression model which attempted to predict college enrollment using both persistence measures and affect and disengagement variables. We chose to use logistic regression because the outcome measure was binary (whether or not the student enrolled in college), and logistic regression is a particularly interpretable algorithm. Several variables came up non-significant (see below). Therefore, in the third analysis, we took the single logistic regression model from the second analysis, removed all non-significant variables, and re-ran the analysis. In these analyses, we again use a Benjamini & Hochberg post-hoc control.

Table 2. Features for Students who attended college (1, n=136) and didn't attend college (0, n=102)

Features (df= 234)	Enrolled in College	Mean	Std. Dev.	t-value	Cohen's D
Mastery Speed	0	**8.327**	**7.916**	**3.816**	0.514
	1	**13.05**	**10.320**	**(p<0.001)**	
Persisted	0	**0.349**	**0.341**	**3.661**	0.489
	1	**0.533**	**0.407**	**(p<0.001)**	
Percent Mastered	0	**0.601**	**0.345**	**4.911**	0.639
	1	**0.802**	**0.278**	**(p<0.001)**	
Quickly Mastered	0	0.458	0.325	-1.831	0.157
	1	0.403	0.369	(p=0.238)	
Persistent-Mastered	**0**	**0.143**	**0.253**	**5.227**	0.715
	1	**0.398**	**0.437**	**(p<0.001)**	
Wheel-Spinning	0	0.136	0.274	-1.261	0.166
	1	0.092	0.253	(p=0.208)	
Quit	**0**	**0.261**	**0.277**	**-5.568**	0.706
	1	**0.105**	**0.148**	**(p<0.001)**	
Percent Correct	**0**	**0.739**	**0.165**	**2.614**	0.342
	1	**0.792**	**0.143**	**(p< 0.001)**	
Boredom	**0**	**0.176**	**0.087**	**-5.3811**	0.719
	1	**0.108**	**0.103**	**(p<0.001)**	
Confusion	**0**	**0.030**	**0.045**	**-5.746**	0.716
	1	**0.006**	**0.016**	**(p<0.001)**	
Concentration	0	0.691	0.078	0.601	0.081
	1	0.698	0.011	(p= 0.547)	
Frustration	0	0.151	0.143	-1.346	0.181
	1	0.121	0.189	(p=0.179)	

Gaming	0	0.139	0.166	-1.3607	0.178
	1	0.111	0.143	(p= 0.174)	
Off Task	0	0.227	0.093	-0.812	0.106
	1	0.218	0.078	(p = 0.418)	

7 Results

7.1 Single-feature analyses

As shown in Table 2, eight features were statistically significant predictors of college enrollment when taken individually, after post-hoc Benjamini & Hochberg correction. Students who enrolled in college had a higher proportion of being persistent in the face of difficulty and mastering difficult topics (Persistent Mastered) (M=0.398) than students who did not enroll in college (M=0.143), t(233)=5.227, p<0.001, D = 0.715. Students who enrolled in college were less likely to quit problem sets without reaching mastery (M=0.261) than students who did not eventually enroll in college (M=0.105), t(233)= -5.568, p<0.001, D = 0.706. However, contrary to our initial hypothesis, wheel-spinning was not significantly different for students who enrolled in college than students who did not, t(233)= -1.261, p=0.208; mastering problem sets quickly was also not statistically significantly different, t(233)=-1.831. In fact, students who enrolled in college actually took longer on average to master the problem sets they mastered (M=13.05) than students who did not enroll (M=8.327), t(233)=3.816, p<0.001, likely due to the difference in quitting early.

There were affective differences between students who enrolled in college versus those who did not enroll, broadly similar to the analyses of the super-set of this data set seen in [28]. Students who enrolled in college were statistically significantly less often bored and confused than students who did not enroll in college. Unlike that larger data set, however, students who enrolled in college did not game the system significantly more or less often than students who did not enroll in college.

7.2 Productive Persistence Models (Full Model)

Table 3 presents a logistic regression model that uses three of the four productive persistence features to predict college enrollment (the fourth, Quickly Mastered, is omitted due to collinearity), controlling for affect and disengaged behavior. We find that productive persistence is still a significant predictor of college enrollment, even after accounting for all the affect and behavior features (p<0.01). However, quitting was no longer a significant predictor of enrollment once other features were controlled for. It is possible that boredom or confusion, both negatively associated with student outcomes, may have played a role both in why students quit and why they are less likely to enroll in college. Engaged concentration was marginally significantly positively associated with enrollment. Gaming the system, on the other hand, became significantly negatively associated with enrollment once other features were controlled for.

Table 3. Logistic regression model including both persistence features and affect and disengagement features. Significant features are in boldface, and marginally significant features are in italics. Alpha values from the Benjamini & Hochberg post-hoc control are included.

Feature	Coefficient	p-value (alpha value)	Odds Ratio	Risk Ratio
Persistent Mastered	**0.350**	**0.004 (0.006)**	**1.420**	**1.606**
Wheel-Spinning	0.308	0.105 (n/a)	1.361	1.678
Quit	-0.108	0.546 (n/a)	0.897	0.569
Boredom	**-0.927**	**0.020 (0.022)**	0.396	0.317
Confusion	**-2.709**	**0.006 (0.011)**	**0.067**	**<0.001**
Engaged Concentration	0.824	0.0279 (0.0278)	2.280	3.829
Frustration	0.164	0.395 (n/a)	1.178	1.390
Gaming	**-0.771**	**0.0096 (0.017)**	**0.463**	**0.263**
Off Task	0.148	0.674 (n/a)	1.159	1.449
Constant	-0.009	0.731 (n/a)	1.104	0.265

Table 4. Model for College Enrollment omitting non-significant predictors. Significant features are in boldface

Feature	Coefficient	p-value (alpha value)	Odds Ratio	Risk Ratio
Persistent Mastered	**0.323**	**0.0005 (0.01)**	**1.382**	**1.663**
Boredom	**-0.814**	**0.0224 (0.03)**	**0.443**	**0.426**
Confusion	**-2.813**	**0.0015 (0.02)**	**0.060**	**0.000**
Engaged Concentration	**0.686**	**0.0453 (0.05)**	**1.986**	**3.170**
Gaming	**-0.408**	**0.0347 (0.04)**	**0.665**	**0.472**
Constant	0.214	0.415 (n/a)	1.239	0.301

Finally, we removed the non-significant features from the model presented in Table 3 in order to examine whether the directionality and significance of the previously significant features remain the same as before. Engaged concentration goes from marginally significant to significant; the significance and directionality of the other features remains the same. (See Table 4)

8 Discussion and Conclusions

In this paper, we compared behaviors associated with persistence in online learning, during middle school, to the longitudinal outcome of whether a student eventually enrolls in college: wheel-spinning (completing many items and never mastering the skill), productive persistence (completing many items and eventually mastering the skill), and quitting a skill (without mastery or completing a substantial number of items).

Of these three, only productive persistence is reliably associated with college enrollment after controlling for student affect and disengaged behaviors. Students who are productively persistent more often are more likely to enroll in college, years later. In fact, students who enroll in college were productively persistent in middle school mathematics almost three times as often, an effect size (Cohen's D) of 0.715, aligning to other research showing the importance of persistence for life outcomes [9, 11, 12].

By contrast, wheel-spinning is not statistically significantly associated with eventual college enrollment. This finding suggests that though struggling without success is an emotionally upsetting experience and is associated with lower amounts of positive engaged concentration [3], it may not be as problematic for students in the long-term as may have been thought. Students may find another way to learn the material they cannot succeed on in the learning system, perhaps asking a teacher, a parent, or fellow students, or perhaps learning it on their own later in the school year.

Quitting without persisting or reaching mastery is associated with statistically significantly lower probability of college enrollment, when taken on its own. However, when controlling for student affect and another form of disengaged behavior, quitting is no longer a statistically significant predictor. This finding suggests an important role for affect and/or disengagement in the processes that mediate between giving up on a problem set and eventual impact on outcomes. It is possible that either boredom or confusion may lead a student to quit a problem set, and the relationship between those two affective states and longer-term outcomes may be a more important factor than how often a student quit a problem set. As with wheel-spinning, students may find another way to learn this material. Better understanding these relationships will be an important area for future work.

Beyond this, this paper's findings suggest several additional directions for future work. Better understanding the intermediate steps between productive persistence in middle school online learning, and eventual outcomes, would likely be a valuable area for further investigation. Is the possible impact of productive persistence due to greater learning? Or is it because productive persistence in online learning correlates to persistence and grit in other contexts as well? Teasing out the degree to which productive persistence is important in itself, and the degree to which it is simply indicative of broader grit, will be a worthwhile question to answer.

At the same time, it may be worth looking further into the future, as has been done for affect [1, 27]. It is not feasible with the current data set – although both skill builder problem sets (needed to assess wheel-spinning and productive persistence according to our current definitions of each) and later longitudinal measures are available for the current data set, the overlap between these variables in insufficiently large to support analysis (i.e. the students who completed skill builders were generally not the same ones for whom later longitudinal measures were available). Investigating this issue will need to wait for another data set.

As a final theme, this paper shows the value of longitudinal data, linked to online learning data, for a variety of secondary analyses. The ASSISTments longitudinal data set has now been used in dozens of analyses beyond its initial intended analyses – at the time that data set was collected, the first paper on wheel-spinning in online learning had not yet been published. All too often, online learning researchers do not retain the information needed for longitudinal follow-up. We hope that this paper provides yet further justification encouraging researchers to retain this information.

References

1. Almeda, M.V.Q. and R.S. Baker, Predicting student participation in STEM careers: The role of affect and engagement during middle school. Journal of Educational Data Mining, 2020. **12**(2): p. 33-47.
2. Arnold, K.E. and M.D. Pistilli. Course signals at Purdue: Using learning analytics to increase student success. in Proceedings of the 2nd international conference on learning analytics and knowledge. 2012.
3. Beck, J. and M.M.T. Rodrigo. Understanding wheel spinning in the context of affective factors. in Proceedings of the International conference on intelligent tutoring systems. 2014. Springer.
4. Beck, J.E. and Y. Gong. Wheel-spinning: Students who fail to master a skill. in Proceedings of the International conference on artificial intelligence in education. 2013. Springer.
5. Benjamini, Y. and Y. Hochberg, Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal statistical society: series B (Methodological), 1995. **57**(1): p. 289-300.
6. Botelho, A.F., Varatharaj, A., Inwegen, E.G.V., and Heffernan, N.T. Refusing to Try: Characterizing Early Stopout on Student Assignments. in Proceedings of the 9th International Conference on Learning Analytics & Knowledge. 2019.
7. Botelho, A.F., Varatharaj, A., Patikron, T., Doherty, D., Adjei, S.A., and Beck, J.E. Developing early detectors of student attrition and wheel spinning using deep learning. IEEE Transactions on Learning Technologies, 2019. **12**(2): p. 158-170.
8. CCSS-MA, Common Core State Standards for Mathematics., 2010: Washington, DC.
9. Credé, M., M.C. Tynan, and P.D. Harms, Much ado about grit: a meta-analytic synthesis of the grit literature. Journal of Personality and social Psychology, 2017. **113**(3): p. 492.
10. Dekker, G.W., M. Pechenizkiy, and J.M. Vleeshouwers, Predicting Students Drop Out: A Case Study. Proceedings of the International Conference on Educational Data Mining, 2009.
11. Duckworth, A., Grit: The power of passion and perseverance. 2016: Scribner New York, NY.
12. Duckworth, A.L., Peterson, C., Matthews, M.D., and Kelly, D.R. Grit: perseverance and passion for long-term goals. Journal of personality and social psychology, 2007. **92**(6): p. 1087.
13. Flores, R.M. and M.M.T. Rodrigo, Wheel-Spinning Models in a Novice Programming Context. Journal of Educational Computing Research, 2020: p. 0735633120906063.
14. Gardner, J. and C. Brooks, Student success prediction in MOOCs. User Modeling and User-Adapted Interaction, 2018. **28**(2): p. 127-203.
15. Gong, Y. and J.E. Beck. Towards detecting wheel-spinning: Future failure in mastery learning. in Proceedings of the second ACM conference on learning@ scale. 2015.
16. Heffernan, N.T. and C.L. Heffernan, The ASSISTments Ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. International Journal of Artificial Intelligence in Education, 2014. **24**(4): p. 470-497.
17. Kai, S., Almeda, M.V., Baker, R.S., Heffernan, C., and Heffernan, N. Decision tree modeling of wheel-spinning and productive persistence in skill builders. Journal of Educational Data Mining, 2018. **10**(1): p. 36-71.

18. Käser, T., S. Klingler, and M. Gross. When to stop? Towards universal instructional policies. in Proceedings of the sixth international conference on learning analytics & knowledge. 2016.
19. Kloft, M., Stiehler, F., Zheng, Z., and Pinkwart, N. Predicting MOOC dropout over weeks using machine learning methods. in Proceedings of the EMNLP 2014 workshop on analysis of large scale social interaction in MOOCs. 2014.
20. Matsuda, N., S. Chandrasekaran, and J.C. Stamper. How quickly can wheel spinning be detected? in Proceedings of the International Conference on Educational Data Mining. 2016.
21. Milliron, M.D., L. Malcolm, and D. Kil, Insight and Action Analytics: Three Case Studies to Consider. Research & Practice in Assessment, 2014. **9**: p. 70-89.
22. Ocumpaugh, J., Baker, R.S., Rodrigo, M.M.T. Baker Rodrigo Ocumpaugh monitoring protocol (BROMP) 2.0 technical and training manual. New York, NY and Manila, Philippines: Teachers College, Columbia University and Ateneo Laboratory for the Learning Sciences, 2015.
23. Ostrow, K., Donnelly, C., Adjei, S., and Heffernan, N. Improving student modeling through partial credit and problem difficulty. in Proceedings of the Second ACM Conference on Learning@ Scale. 2015. ACM.
24. Pardos, Z.A., Baker, R.S., San Pedro, M.O., Gowda, S.M., and Gowda, S.M. Affective states and state tests: investigating how affect and engagement during the school year predict end-of-year learning outcomes. Journal of Learning Analytics, 2014. **1**(1): p. 107-128.
25. Patikorn, T., R.S. Baker, and N.T. Heffernan, ASSISTments Longitudinal Data Mining Competition Special Issue: A Preface. Journal of Educational Data Mining, 2020. **12**(2): p. i-xi.
26. Razzaq, L. and N.T. Heffernan. Scaffolding vs. Hints in the Assistment System. Proceedings of the International Conference on Intelligent Tutoring Systems. 2006. Heidelberg, Germany: Springer.
27. San Pedro, M.O., Ocumpaugh, J., Baker, R.S., and Heffernan, N.T. Predicting STEM and Non-STEM College Major Enrollment from Middle School Interaction with Mathematics Educational Software. in Proceedings of the International Conference on Educational Data Mining. 2014.
28. San Pedro, M.O.Z., Baker, R., Bowers, A., and Heffernan, N. Predicting College Enrollment from Student Interaction with an Intelligent Tutoring System in Middle School. in Proceedings of the 6th International Conference on Educational Data Mining. 2013.
29. Tinto, V., Leaving college: Rethinking the causes and cures of student attrition. 1987. Chicago, USA: University of Chicago Press.
30. Van Inwegen, E.G., Adjei, S.A., Wang, Y., and Heffernan, N.T. Using Partial Credit and Response History to Model User Knowledge. in Proceedings of the International Conference on Educational Data Mining, 2015.
31. Wang, Y. and R. Baker, Grit and intention: Why do learners complete MOOCs? The International Review of Research in Open and Distributed Learning, 2018. **19**(3).
32. Wang, Y., N.T. Heffernan, and J.E. Beck. Representing student performance with partial credit. in Proceedings of the International Conference on Educational Data Mining 2010. 2010.
33. Whitehill, J., Williams, J., Lopez, G., Coleman, C., and Reich, J. Beyond prediction: First steps toward automatic intervention in MOOC student stopout. in Proceedings of the International Conference on Educational Data Mining, 2015.
34. Xiong, X., S. Li, and J.E. Beck. Will You Get It Right Next Week: Predict Delayed Performance in Enhanced ITS Mastery Cycle. in Proceedings of the Florida Artificial Intelligence Research Symposium. 2013.

35. Xiong, Y., Li, H., Kornhaber, M.L., Suen, H.K., Pursel, B., and Goins, D.D. Examining the relations among student motivation, engagement, and retention in a MOOC: A structural equation modeling approach. Global Education Review, 2015. **2**(3): p. 23-33.