Research Article

Till Massing*, Natalie Reckmann, Alexander Blasberg, Benjamin Otto#, Christoph Hanck, Michael Goedicke

When is the Best Time to Learn? – Evidence from an Introductory Statistics Course

<https://doi.org/10.1515/edu-2020-0144>

received December 16, 2020; accepted February 23, 2021.

Abstract: We analyze learning data of an e-assessment platform for an introductory mathematical statistics course, more specifically the time of the day when students learn and the time they spend with exercises. We propose statistical models to predict students' success and to describe their behavior with a special focus on the following aspects. First, we find that learning during daytime and not at nighttime is a relevant variable for predicting success in final exams. Second, we observe that good and very good students tend to learn in the afternoon, while some students who failed our course were more likely to study at night but not successfully so. Third, we discuss the average time spent on exercises. Regarding this, students who participated in an exam spent more time doing exercises than students who dropped the course before.

Keywords: Learning analytics, E-assessment, effective learning times

1 Introduction

Modern university courses often use e-assessment systems. Especially when courses have a high number of participants e-learning tools are very useful to give students individual feedback. Courses with quantitative

contents such as statistics and introductory mathematics are particularly suitable for e-assessment since exercises can also be designed as fill-in exercises – which require students to submit a numeric answer – instead of only multiple choice exercises. The e-assessment system JACK is a framework for delivering and grading complex exercises of various kinds (Schwinning et al., 2017). It was originally created to check programming exercises in Java (Striwe et al., 2009), but has been extended to several other exercise types such as multiple-choice and fill-in exercises (Striwe et al., 2015; Striwe, 2016; Schwinning et al., 2017). JACK offers parameterizable content, meaning that exercises can contain different values each time an exercise is practiced. This means not only that different students get a different parameterization but, moreover, that the same student sees different numbers at each different time s/he selects the exercise. Hence, the exercise remains challenging until s/he understands the underlying concept to solve the exercise.

In addition to fill-in exercises, JACK allows to design exercises with dynamic programming content. For instance, JACK offers electronic Java or R – the standard statistical programming language – exercises. Programming exercises not only help to prepare students for modern statistical work, but also have been shown to be highly beneficial to foster their understanding of statistics (Otto et al., 2017; Massing et al., 2018a).

This study analyzes JACK data to more deeply understand students' learning behavior in an introductory mathematical statistics course. The high correlation between learning effort in the semester and the final grades is well documented (Massing et al., 2018b, and Section 2 provide examples). Here, we aim to investigate additional aspects, namely the relevance of the daytime when students learn. More specifically, our research question is to analyze if it is possible to identify recommendations for students on when to best schedule their learning activities. To this end, we exploit granular data of the time stamps of student learning data made available by JACK. In order to do so, we use several statistical learning methods to study

*Corresponding author: Till Massing, Faculty of Economics, University of Duisburg-Essen, Universitätsstr. 12, 45117 Essen, Germany, E-mail: till.massing@uni-due.de

Natalie Reckmann, Alexander Blasberg, Christoph Hanck, Faculty of Economics, University of Duisburg-Essen, Universitätsstr. 12, 45117 Essen, Germany

Benjamin Otto, Michael Goedicke, paluno - The Ruhr Institute for Software Technology, University of Duisburg-Essen, Gerlingstr. 16, 45127 Essen, Germany

For more information regarding JACK contact this author.

which factors in students' learning behavior are relevant to predict their success in the exam (e.g., using tree-based methods as in Hu et al., 2014). It turns out that daytime activity has a higher effectiveness than nighttime activity, which is in line with Gomes et al. (2011).

Additionally, we find that good and very good students favor to learn in the afternoon, while some students who failed our course had insufficient learning behavior late at night. Moreover, we discuss the average time spent on exercises. Regarding this, students who participated at an exam spent more time on exercises than students who dropped the course before. This strengthens findings of Xiong et al. (2011), who also found that response time is a relevant predictor for student performance.

The remainder of this paper is organized as follows: Section 2 provides a brief overview of related work. Section 3 introduces the statistics course analyzed here and Section 3.1 presents the available data and the models used. Section 4 discusses the empirical results. Section 5 concludes.

2 Related work

The overall engagement of students is indisputably one of the main covariates of academic success. For the case of mathematical statistics this has been shown on several occasions. Sosa et al. (2011) show in a meta-study that the simultaneous usage of traditional classroom lectures and e-assessment has a positive effect on students' success. Massing et al. (2018b) substantiate the previous result by analyzing the learning activity on the e-assessment platform JACK. The study reveals that learning effort and success, measured by the total number of (correct) submissions on JACK over the course, positively affects the final grade in the exam. Otto et al. (2017) add additional R-programming exercises to the JACK framework and show that this exercise type helps to improve the general understanding of fundamental statistical concepts and thus ultimately yields better results in the final exam.

Due to the empirically observed positive effect of a multitude of variables on academic performance, prediction of the latter has become possible. In this so-called branch of educational data mining various statistical learning methods are applied to educational data in order to predict student outcomes. Often, but not necessarily, this outcome is measured with a binary response of pass/fail in order to be able to provide an early-warning to students. James et al. (2013); Hastie et al. (2009); Meier et al. (2016) give a comprehensive overview of popular statistical learning methods used in

the literature. For an overview of how to implement an early-warningsystem see, e.g., bin Mat et al. (2013). The literature has identified a number of important predictors. Gray et al. (2014) find evidence for the importance of socioeconomic and psychometric variables as well as pre-university grades, although Oskoue and Askari (2014) show that, especially among the socioeconomic variables, the predictive capability can vary across countries. Baars et al. (2017) additionally identify post-admission variables like obtained credits, degree of exam participation and exam success rate to have an influence on students' success. Macfadyen and Dawson (2010); Wolff et al. (2013); Elbadrawy et al. (2015) analyze the learning activity on learning management systems and are able to accurately predict students' performance with appropriate variables. In a similar but more assessmentbased fashion Huang and Fang (2013); Burgos et al. (2018); Massing et al. (2018a) use activity in e-learning frameworks as well as the results of mid-term exams to predict students' success in the final exam. Asif et al. (2017) identify the performance in a small number of selected courses as a predictor for the academic achievement at the end of the study program. In contrast to previous studies which mostly rely on quantitative and qualitative learning activity measured by time-invariant variables, there is also a temporal dimension of engagement which has been studied from different perspectives throughout the academic literature. Hu et al. (2014) use time-dependent information provided by a learning management system to predict academic performance. Xiong et al. (2011) incorporate students' response time as an additional feature into a random forest to investigate the predictive capability for students' performance and find evidence that it can indeed improve prediction accuracy. Papamitsiou et al. (2014) and Papamitsiou et al. (2016) further elaborate on the latter by using more sophisticated techniques and are able to support the preceding result.

Only a few studies focus on the intraday engagement of students, that is, the actual time of the day of learning, as a predictor for academic success. This topic is relevant as various studies show a significant influence of sleep quality and patterns on academic performance (Ahrberg et al., 2012; Baert et al., 2015; Beşoluk et al., 2011; Eliasson et al., 2002). Based on these insights Gomes et al. (2011) incorporate sleep variables into a prediction setting. With a stepwise regression approach they identify sleep frequency, night outings and sleep quality as among the most important predictors of academic success.

3 Course Structure

This section outlines the initial setup of the study. In particular, we sketch the structure of the analyzed course.

The e-assessment system JACK was used for a lecture and exercises course in mathematical statistics at the German university of Duisburg-Essen. 753 undergraduate first-year students started the course. The course is compulsory for several business and economics programs as well as in teachers' education. Out of these 753 students, only 379 took an exam at the end of the course, while the others dropped the course in this term (Table 1). The course also introduces statistical programming skills using the statistics software R. In order to do so, the e-assessment system JACK offers programming exercises where the correctness of students' code is assessed, in addition to classical fill-in and multiple-choice exercises.

The course consisted of a weekly 2-hours lecture, which introduced statistical concepts, and a 2-hours exercise class, which presented explanatory exercises and problems. Both classes were held classically in front of the auditorium. Due to the large number of students, these classes are limited in addressing students' different speeds of learning and individual questions. To overcome this issue and to encourage self-reliant learning, as well as to support students who had difficulties to attend classes, we offered all homework on JACK.

All in all, we offer 173 different exercises on JACK, of which 48 are designed as R-programming exercises and the remainder as fill-in or multiple-choice exercises. The individual learning success is supported by offering specific automated feedback and, furthermore, by optional hints. In case of additional questions which were neither covered by hints nor feedback, the students were able to ask questions in our moodle help-forum.

In order to further encourage students to learn continuously during the semester, and not only in the weeks prior to the exams, we offered five online tests using JACK. These tests lasted 40 minutes at fixed times in the evening. Four of the online tests contained fill-in or multiple-choice exercises only. The fifth online test contained R exercises exclusively. Participation only required a device with internet access, but no compulsory attendance at university. This summative assessment allowed students to assess their individual state of knowledge during the lecture period. It was not compulsory for students to participate at online tests in order to take the final exam at the end of the course. Instead, we offered bonus points for the final exams to encourage participating at the tests (a maximum of 10

Table 1: Overview of the number of students registered to the course and the number of submissions on JACK.

Students, who. . .	counts	#homework submissions
took course	753	163,444
registered to an exam	438	152,232
participated at an exam	379	147,868
passed an exam	186	87,382

bonus points in total for fill-in online tests). The bonus points were only added to final exam points if at least 25 out of 60 exam points were reached, i.e., if students passed the exam without bonus. The R online test was worth at most 2 bonus points which were awarded even if students achieved less than 25 points. The reason for this was to motivate students to focus on programming skills since Otto et al. (2017) and Massing et al. (2018a) show that this has a substantially (three times) higher impact on exam success than classical fill-in exercises.

The final exams (3 in total) were also held electronically. While online tests during the semester could be solved at home with open books, the final exams were offered exclusively at university PC pools and supervised by academic staff. The exam consisted of R exercises (15%), short handwritten proofs (15%) and the remainder of fill-in exercises. Students can only retake an exam if they failed or did not take the previous ones (so that students can pass at most once), but can fail several times.¹ The last grade a student achieved in an exam will be denoted as the final grade. The corresponding exam will be denoted as the final exam.

3.1 Data and Models

In this section we present the available database and the models used. For each homework submission by a student on JACK we observe the exercise ID, the student ID, the number of points (on a scale from 0 to 100) and the time stamp with minute-long precision.

The response variables are given by the final exam success. We consider two possible responses. The first is a binary variable indicating whether a student passed (1)

¹ Students obtain 6 “malus points” for each failed exam of which they may collect at most 180 during their whole bachelor program.

or did not pass (0) the course. Second, we consider the final grade as a response. We have the following grading scheme: very good (“100”), good (“200”), satisfactory (“300”), sufficient to pass (“400”) and failed (“500”). We assign “600” to students who took the course but did not participate at any of the exams. This is actually not a grade. However, this reflects the view that students who did not take any exam were even less prepared than students who failed the exams. Table 2 reports an overview of final grades. We do not report grades for one specific exam date but the grade given at the end of the course.

JACK registered 163,444 submissions of homework exercises. See Table 1 for how these submissions are distributed among students. Figure 1 plots the number of daily submissions on JACK aggregated for all students. Characteristically, the number of submissions peaks shortly before a summative assessment such as an exam or an online test. This was also observed by Massing et al. (2018b).

We compile the following information for each student i from the raw data:

- the number of submissions (# submissions in short),
- the number of fully correct submissions (100 points),
- the number of submissions in the morning from 8am to 12pm,
- the number of submissions in the afternoon from 12pm to 4pm,
- the number of submissions in the evening from 4pm to 8pm,
- the number of submissions in the late evening from 8pm to 12am,
- the number of submissions at night from 12am to 8am,
- the median submission time (Subsection 4.2).
- The score, which is defined as follows: let t be a day during the semester. Then

$$score_{it} := \sum_{j=1}^n x_{ijt},$$

where x_{ijt} is the number of points of the latest submission up to time t of student i in exercise j , $j = 1, \dots, n$. In other words, the score is the sum of points of the last submissions to every exercise. This helps tracking the learning progress for every student. In particular, we consider the final score, which is the score evaluated at the end of the term.

- The frequency of submissions, i.e., the mean time between two following submissions at different days, measured in days.
- The time until a student hands in the first submission from the beginning of the term, measured in days.

Table 2: Overview of the distribution of grades, with the failure rate in the second-to-last line. The last line gives the “average grade”.

Grade (points)	counts
“100” (48 – 60)	9
“200” (40 – 47)	38
“300” (31 – 39)	100
“400” (25 – 30)	39
“500” (0 – 24)	193
Σ	379
failure rate	.508
“average grade”	396.4
“standard deviation”	117.8

Figure 1: Number of daily submissions on JACK versus the date. Exam and online test dates are highlighted with different points.

- The time until a student hands in the last submission before his/her last exam, measured in days.
- The number of days a student submitted solutions.
- The average time spent per exercise measured in minutes (Subsection 4.3).

Table 3 reports summary statistics. Figure 2 plots the average score of students with different grades and of students who dropped the course. Evidently, good and very good students had a strong learning progress from the beginning of the semester on. Students with the sufficient pass grade “400” and students who failed

Table 3: Overview of empirical quartiles mean and standard deviation for the considered covariates.

Variable name	Min	1st Quartile	Median	Mean	3rd Quartile	Max	Sd
# submissions	0	16	85	219.2	329	2726	296.8
# correct	0	7	32	94.7	126	1576	144.3
# morning	0	0	9	32.97	38	653	62.2
# afternoon	0	3	28	73.75	103	1273	111.7
# evening	0	5	32	71.56	105	819	97.9
# late ev	0	0	5	31.52	36	465	58.3
# night	0	0	0	9.432	5	396	32
Median submission time	12:17am	2:16pm	3:56pm	3:37pm	7:19pm	11:31pm	2:52
Final score	0	500	2372	4105	7333	15681	4100
Frequency	1	5.25	9.556	50.433	26	234	84.8
First submission	0	7	25.3	36.69	29.3	259.3	48.7
Last submission	0	0.5	17.9	75.71	173.5	234	88.1
# days	0	2	9	15.47	25	90	16.6
Ave Time spent	0	5.8	10.3	10.5	13.9	46.3	6.8

(“500”) start similarly weak but improve shortly before the exam. Students with “400”, however, improve slightly more, which may be the reason that they pass the exam. On the other hand, they may just have been lucky in the exam. The students who dropped the course show very little progress on average.

We choose the following modeling approaches for the classification problem:²

- *Logistic Regression.* Logistic regression models the probability of an event given p regressor variables $X = (X_1, \dots, X_p)$ via

$$p(X) = \frac{\exp(\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p)}{1 + \exp(\beta_0 + \beta_1 X_1 + \dots + \beta_p X_p)}. \quad (1)$$

The idea is to regress the log-odds, $\log\left(\frac{p(X)}{1-p(X)}\right)$, on a linear combination of X . So equation (1) can be rewritten as

$$\log\left(\frac{p(X)}{1-p(X)}\right) = \beta_0 + \beta_1 X_1 + \dots + \beta_p X_p. \quad (2)$$

² We also tried other modeling approaches than the ones stated here (e.g. neural networks, support vector machines, etc.). However their predictive performance proved not to be competitive.

- The unknown coefficients $\beta_0, \beta_1, \dots, \beta_p$ are estimated based on the available data. We use maximum likelihood (James et al., 2013). We measure the variability of the estimates $\hat{\beta}_0, \hat{\beta}_1, \dots, \hat{\beta}_p$ via the standard errors $SE(\hat{\beta}_l), l = 0, 1, \dots, p$, of the estimates. From this we obtain the t -statistics $t_l = \frac{\hat{\beta}_l}{SE(\hat{\beta}_l)}$.

. Kuhn (2012) recommends using the absolute value of the t -statistic of each non-constant regressor as an importance measure for logistic regression.

The above approach can easily be extended to ordered logistic regression in which we want to predict a variable with $k > 2$ possible outcomes (multi-class-classification) (Agresti, 2002). We use binary logistic regression to predict the response “student passed”, and ordered logistic regression to predict the grade.

- *Random Forests.* Tree-based methods can be used for two-class-classification as well as multiclass-classification. Single decision trees are very easy to interpret but have the drawback of having a high variance. To avoid this problem Breiman (2001) proposed an algorithm for averaging decision trees to obtain a so-called random forest. The idea is to take B bootstrap samples from the single training data set. Then, a tree is trained on every bootstrapped training data sample. Finally, the prediction is the majority vote, which is the most common occurring class over all B predictions (James et al., 2013). Each of the single trees has a high variance but a low bias. Averaging

over all trees reduces the variance. Another problem is that in each split of the trees, every variable in the predictor space is considered. If there is, for example, one very strong predictor it will be used in each tree for the first split. This leads to a high correlation between the trees. To avoid this problem, only a random sample of the p predictors is used in each split to find an optimal split. The number of predictors in this random sample is usually set to $\sqrt{p}$. Breiman (2001) also proposed to use the *Mean Decrease Accuracy* as importance measure for the input variables. We build 500 trees to grow the forest and try 3 variables at each split.

4 Empirical results

This section analyzes students' learning behavior. We discuss which learning strategy turns out to predict students' exam success.

4.1 Variable Importance

Our first analysis discusses which of the explanatory variables have a high predictive relevance. To model the target variable, i.e., passing the course or achieving a certain grade, we use the following set of variables as predictors in all of our models: # correct, # morning, # afternoon, # evening, # late evening, # night, score at day of first online test, score at day of third online test, final score, frequency, first submission, last submission, # days, total time spent for exercises. We dropped some variables like the total number of submissions to avoid high correlation between the predictor variables.

To compare the performance of the two models we use the accuracy, i.e., the rate of correctly classified observations. To avoid overfitting we use 3-fold cross validation. Table 4 contains the cross validation results for the two different models. We see that in both, two-class and multi-class classification, the random forest works best with an accuracy of 0.830 or 0.73 but logistic regression works well, too. In the full data set 75.3% of the students do not pass the course, which leads to an accuracy of 0.753 if we predict all students to not pass the course. Hence the random forest leads to an increase in accuracy of around 8 percentage points which leads us to use the results of the random forest from now on.

Table 4: Accuracy of both models.

Model	Accuracy two-class	Accuracy multi-class
Logistic Regression	0.821	0.72
Random Forest	0.830	0.73

Figure 2: Average score (averaged over students with same grade) versus the date. Thin vertical lines mark the exam dates.

We now investigate the variables which are chosen to build the single trees for the random forest. Figure 3 shows the importance of the variables used in the analysis.

We can see that the variable *last submission*, i.e., the time until a student hands in the last submission before his/her last exam, measured in days, is by far the most important variable. Unfortunately, Figure 3 is silent on the direction of impact on the target variable. A solution to this problem is the partial dependence plot, which can help to understand how the log-odds of realizing the respective class depend on the input variables.³ A high positive value of the partial dependence for a given value of the predictor means that it is more likely to belong to the class of interest than to the other class, see Hastie et al. (2009). Here the class of interest is *not passing* the course. Figure 4 shows the partial dependence plot for *last submission*. We see that not passing the course is more likely for high values of *last submission*. This means

³ The y-axis shows $f(x) = \log p_k(x) - \frac{1}{K} \sum_{j=1}^K \log p_j(x)$, where K is the number of classes, k is the class of interest, and p_j is the proportion of votes for class j .

Figure 3: Variable importance measured in mean decrease accuracy for the two-class random forest.

that students who learn until the day of the final exam unsurprisingly have a higher probability to pass the course than students who quit learning far before the exam. This is because 374 out of 753 students did not participate in an exam. Most of these students did not learn until the exam but only made a few submissions at the beginning of the semester. Hence the variable *last submission* has high values for these students. On the other hand most of the students who participated in the exams learned until shortly before the exam. This implies the high importance of the last submission. Other important variables are the final score, the frequency of submissions and the number of submissions in the morning. Figure 5 shows the partial dependence plot of the final score. We see that a high final score leads to a low probability not to pass the course.⁴

Furthermore, the importance of the variables in Figure 3 shows that the time of the first submission, the number of submissions at night and in the late evening and the first score in the term do not help for the predictive performance in the final exam. For the former and the latter this could be due to the fact that, at the beginning of the course, almost all students start to learn at the same level of knowledge, so there is no information that helps to decide between students passing or failing the final exam.

⁴ Note that in logistic regression the sign of the estimated coefficients tells the direction of the impact of a variable. These are mostly in line with the exemplary partial dependence plots.

Figure 4: Partial dependence plot for the last sub-mission variable before the exam for the multi-class random forest. Vertical ticks on the x-axis indicate deciles of the last submission variable.

Since 374 out of 753 students did not participate in the exams we only focus on students who participated in an exam for the remainder of this subsection. This will obviously reduce the impact of the variable *last submission*. We now estimate the corresponding binary classification random forest for pass vs. fail. Figure 6 shows its variable importance plot.⁵ Now, final score and frequency, i.e., the mean time between two days of submissions measured in days, are the most important variables in the random forest model. For example, Figure 7 shows the partial dependence plot for the frequency. Small values of frequency make it more likely to pass the course. This means that students who learn regularly with only a few hours between their submissions have a higher probability of passing.

In case of multi-class classification, again, the time until a student hands in the last submission before his/her last exam is by far the most important variable in the model, for the same reasons as above. All other variables have low importance in this model. For reasons of brevity we shall now focus on the results of the binary model.

4.2 Learning Times

We now analyze more deeply at which time of the day good and less successful students prefer to learn. In order to

⁵ Note that a negative value for the mean decrease accuracy implies that randomly permuting the respective variable (*ceteris paribus*) yields to a lower MSE of the random forest.

Figure 5: Partial dependence plot for the final score variable for the multi-class random forest. Vertical ticks on the x-axis indicate deciles of the final score variable.

Figure 6: Variable importance measured in mean decrease accuracy for the two-class random forest without students who dropped the course.

investigate this we compute the median submission time for each student. We compare the median submission times for students who passed or did not pass in final exams.

Figure 8 shows kernel density plots for the median submission time for passing students in solid black and

non-passing students in dashed red. There is a higher variance of median submission times for students who did not pass; students who passed prefer to learn in the afternoon. Weaker students tend to learn later. Moreover, quite a few non-passing students have median learning time in the morning. This is usually the time of the day when students should attend lecture and exercise classes.

Figure 9 further supports this claim. We compare very good and good students with students who failed all exams and students who dropped the course. Evidently, good students prefer to submit exercises during daytime. The earliest median submission time of a good student is about 11:30am and the latest is 7:20pm. Comparing these students with students who failed and, more visibly, who dropped shows that there are quite a few who study very late or very early. For example, the earliest is about 12:20am and the latest about 11:30pm.

This leads us to conclude that there are more nonpassing students who have difficulties in learning in the afternoon. As stated in Section 2 lack of sleep caused by studying at night has a negative impact on students' performance.

Needless to say, the more important reason that poor students fail is mainly because they learn too little and not because of bad timing (cf. Subsection 4.1). It also needs to be emphasized that unfavorable time management can also be due to a high amount of responsibilities not connected to their studies. Unfortunately, our data set does not allow us to distinguish between these aspects. A data set including both submission data for e-assessment and information on students' other daily activities is hard to collect.

4.3 Submission duration

We now highlight another influential factor for success: we consider how long students work on a single submission, i.e., how much time they spent solving an exercise. This analysis faces some challenges. First, we only observe the end and not the beginning of solving an exercise and hence do not have exact start and end times. We bypass this problem by measuring the time between two successive submissions. For example, if a student submits an exercise at 12pm and submits a second exercise (which might be the same as the first) at 12:15pm we consider 15 minutes as time spent for the second exercise. This means we do not observe duration of the first submission but of the following submissions. We omit duration times which are longer than two hours because students then likely took a break. This is also part of the second issue because

Figure 7: Partial dependence plot for the frequency variable for the two-class random forest without students who dropped the course. Vertical ticks on the x-axis indicate deciles of the frequency variable.

Figure 9: Kernel density estimates of the median daytime of submissions on JACK comparing students with different grades and with students who dropped the course.

Figure 8: Kernel density estimates of the median daytime of submissions on JACK comparing students who passed with students who did not pass.

we only monitor submission times in JACK and not whether students used this time to learn or whether they got distracted. We cannot rule out times of distractedness but still believe the following analysis offers interesting insights.

For each student we accumulate all duration times. These totals are of course higher for students who

submitted more exercises than for students who submitted only a few. We thus divide total duration by the number of exercises submitted for each student.

Figure 10 shows a kernel density estimate for the average time spent per submission of students who passed *versus* students who did not pass. Evidently, there are many students who did not pass who invested little time for each exercise. Again, we next distinguish between students who failed an exam and students who dropped the course. Figure 11 compares students who achieved the best or second best grade with students who failed and students who dropped the course. Interestingly, time spent per submission is similar for both good students and students who failed (The plot for mediocre students looks very similar, too). However, students who dropped the course invested perceptibly less time for each submission. Apparently, these students had too little motivation and/or time to participate in the course. They likely did not seriously attempt to solve the exercises.

5 Conclusions and Discussion

This study analyzes when students should learn to be successful in a final exam. For this purpose, we analyzed data from the online learning platform JACK from an introductory mathematical statistics course in the summer term 2017. This data on students' submissions on JACK

Figure 10: Kernel density estimates of the average time-spent for submissions comparing students who passed with students who did not pass.

Figure 11: Kernel density estimates of the average time-spent for submissions comparing students with different grades and with students who dropped the course.

offered information about the daytime when a student submits a solution to an exercise.

We used logistic regression and random forests to predict the success of a student in the final exam and, also try to predict the final grade. An advantage of these methods is that they offer information about the

importance of the variables used in the model. We analyze the variable importance obtained by the random forest.

The two most important variables in this model are the day between the last submission of an exercise and the exam as well as the score the students achieve when they study with JACK. We further identify the frequency with which the students work on JACK and the number of submissions between 8am and 12pm as important variables. We identify good students to submit exercises during the daytime, while some students who quit the course or fail in the final exam learn very early in the morning or very late in the evening. Needless to say, the total amount of learning has a high impact on success. Additionally, we cannot rule out external factors (e.g. working during daytime) causing this effect rather than students who purposely did not study during daytime. Still we may conclude that students who did not pass the course study little during the afternoon. Moreover the time a student spends on a single exercise is very short for students who dropped the course.

All in all, our results strengthen the evidence of the relevance of when and how often students learn. Suitable time management may increase students' probability to pass a course like the one investigated here. That said, our study has some limitations which suggest potential for future work. First, we do not observe the reason why students learn at nighttime. However, these underlying reasons, such as childcare tasks or part-time work, may be the actual drivers of poor exam performance. Hence, the positive association between nighttime studying and poor performance may not so much be “bad” time management, but rather a conscious consequence of – at least some – students' priorities that lead to lower academic success.

To estimate an actual causal effect of bad time management we would require such and possibly further information. One direction for future research is to conduct conditional surveys on JACK to obtain such relevant identifying information. A second limitation is that offline learning time, e.g., working with the lectures slides or textbooks is not covered in our setting since we cannot monitor such efforts electronically. However, we are confident that this aspect is not of dominant importance as high submission numbers suggest that students focus their efforts on JACK (which is also confirmed by informal student feedback). Third, as of now, JACK does not monitor “distractions” of students as, e.g., long time of passivity when working on an exercise (e.g., switching to other browser tabs, watching TV). Hence, a long measured duration of an attempt to solve an exercise may in fact not so much reflect sustained effort, but rather be due to a

distraction. Future work using technical (such as mouse tracking) or survey-based approaches shall try to address these issues.

Acknowledgments: We thank all colleagues who contributed to the course “Induktive Statistik” in the summer term 2017, especially Paul Navas Alban. We thank Anna Janssen, Kim J. Hermann, Janine Langerbein, Timo Rammert and Alexander Zyweck for excellent research assistance.

We gratefully acknowledge the valuable advice by two anonymous reviewers which helped to substantially improve this paper.

Funding Information: Part of the work on this project was funded by the German Federal Ministry of Education and Research under grant numbers 01PL11075 and 01JA 1610.

Conflict of interest: The authors state no conflict of interest.

References

- Agresti, A. (2002). *Categorical Data Analysis*. Hoboken, New Jersey: John Wiley & Sons Inc. <https://doi.org/10.1002/0471249688>
- Ahrberg, K., Dresler, M., Niedermaier, S., Steiger, A., & Genzel, L. (2012). The interaction between sleep quality and academic performance. *Journal of Psychiatric Research*, 46(12), 1618–1622. <https://doi.org/10.1016/j.jpsychires.2012.09.008>
- Asif, R., Merceron, A., Ali, S. A., & Haider, N. G. (2017). Analyzing undergraduate students' performance using educational data mining. *Computers & Education*, 113, 177–194. <https://doi.org/10.1016/j.compedu.2017.05.007>
- Baars, G. J. A., Stijnen, T., & Splinter, T. A. W. (2017). A model to predict student failure in the first year of the undergraduate medical curriculum. *Health Profession Education*, 3(1), 5–14. <https://doi.org/10.1016/j.hpe.2017.01.001>
- Baert, S., Omeij, E., Verhaest, D., & Vermeir, A. (2015). Mister Sandman, bring me good marks! On the relationship between sleep quality and academic achievement. *Social Science & Medicine*, 130, 91–98. <https://doi.org/10.1016/j.socscimed.2015.02.011>
- Beşoluk, S., Önder, İ., & Deveci, İ. (2011). Morningness-eveningness preferences and academic achievement of university students. *Chronobiology International*, 28(2), 118–125. <https://doi.org/10.3109/07420528.2010.540729>
- bin Mat, U., Buniyamin, N., Arsad, P. M., & Kassim, R. (2013). An overview of using academic analytics to predict and improve students' achievement: A proposed proactive intelligent intervention. *5th Conference on Engineering Education (ICEED)*, 126–130. IEEE. doi: 10.1109/ICEED.2013.6908316
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Burgos, C., Campanario, M. L., de la Peña, D., Lara, J. A., Lizcano, D., & Martínez, M. A. (2018). Data mining for modeling students' performance: A tutoring action plan to prevent academic dropout. *Computers & Electrical Engineering*, 66, 541–556. <https://doi.org/10.1016/j.compeleceng.2017.03.005>
- Elbadrawy, A., Studham, R. S., & Karypis, G. (2015). Collaborative Multi-Regression Models for Predicting Students' Performance in Course Activities. *Proceedings of the Fifth International Conference on Learning Analytics and Knowledge, LAK '15*, 103–107. ACM. <https://doi.org/10.1145/2723576.2723590>
- Eliasson, A., Eliasson, A., King, J., Gould, B., & Eliasson, A. (2002). Association of sleep and academic performance. *Sleep and Breathing*, 6(1), 45–48. <https://doi.org/10.1055/s-2002-23157>
- Gomes, A. A., Tavares, J., & de Azevedo, M. H. P. (2011). Sleep and academic performance in undergraduates: A multi-measure, multi-predictor approach. *Chronobiology International*, 28(9), 786–801. <https://doi.org/10.3109/07420528.2011.606518>
- Gray, G., McGuinness, C., & Owende, P. (2014). An application of classification models to predict learner progression in tertiary education. *International Advance Computing Conference (IACC)*, 549–554. IEEE. <https://doi.org/10.1109/IADCC.2014.6779384>
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning*. New York: Springer. <https://doi.org/10.1007/978-0-387-84858-7>
- Hu, Y.-H., Lo, C.-L., & Shih, S.-P. (2014). Developing early warning systems to predict students' online learning performance. *Computers in Human Behavior*, 36, 469–478. <https://doi.org/10.1016/j.chb.2014.04.002>
- Huang, S., & Fang, N. (2013). Predicting student academic performance in an engineering dynamics course: A comparison of four types of predictive mathematical models. *Computers & Education*, 61, 133–145. <https://doi.org/10.1016/j.compedu.2012.08.015>
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An Introduction to Statistical Learning: with Applications in R*. New York: Springer. <https://doi.org/10.1007/978-1-4614-7138-7>
- Kuhn, M. (2012). Variable selection using the caret package. URL <http://cran.cermin.lipi.go.id/web/packages/caret/vignettes/caretSelection.pdf>
- Macfadyen, L. P., & Dawson, S. (2010). Mining LMS data to develop an “early warning system” for educators: A proof of concept. *Computers & Education*, 54(2), 588–599. <https://doi.org/10.1016/j.compedu.2009.09.008>
- Massing, T., Reckmann, N., Otto, B., Hermann, K. J., Hanck, C., & Goedicke, M. (2018a). Klausurprognose mit Hilfe von E-Assessment-Nutzerdaten. In *DeLFI 2018 Die 16* (pp. 171–176). E-Learning Fachtagung Informatik. Bonn: Gesellschaft für Informatik e.V.
- Massing, T., Schwinning, N., Striewe, M., Hanck, C., & Goedicke, M. (2018b). E-assessment using variable content exercises in mathematical statistics. *Journal of Statistics Education: An International Journal on the Teaching and Learning of Statistics*, 26(3), 174–189. <https://doi.org/10.1080/10691898.2018.1518121>
- Meier, Y., Xu, J., Atan, O., & van der Schaar, M. (2016). Predicting grades. *IEEE Transactions on Signal Processing*, 64(4), 959–972. <https://doi.org/10.1109/TSP.2015.2496278>
- Oskouei, R. J., & Askari, M. (2014). Predicting Academic Performance with Applying Data Mining Techniques (Generalizing the results of Two Different Case Studies). *Computer Engineering and Applications Journal*, 3(2), 79–88. <https://doi.org/10.18495/comengapp.v3i2.81>

- Otto, B., Massing, T., Schwinning, N., Reckmann, N., Blasberg, A., Schumann, S., Hanck, C., & Goedicke, M. (2017). Evaluation einer Statistiklehrveranstaltung mit dem JACK R-Modul. *DeLFI 2017 Die 15. e-Learning Fachtagung Informatik, Lecture Notes in Informatic, Gesellschaft für Informatik*, 75–86.
- Papamitsiou, Z., Karapistoli, E. & Economides, A. A. (2016). Applying classification techniques on temporal trace data for shaping student behavior models. *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge, LAK '16*, 299–303, ACM. <https://doi.org/10.1145/2883851.2883926>
- Papamitsiou, Z. K., Terzis, V., & Economides, A. A. (2014). Temporal learning analytics for computer based testing. *Proceedings of the Fourth International Conference on Learning Analytics And Knowledge, LAK '14*, 31–35, ACM.
- Schwinning, N., Striwe, M., Massing, T., Hanck, C., & Goedicke, M. (2017). Towards digitalisation of summative and formative assessments in academic teaching of statistics. *Proceedings of the Fifth International Conference on Learning and Teaching in Computing and Engineering*.
- Sosa, G. W., Berger, D. E., Saw, A. T., & Mary, J. C. (2011). Effectiveness of computer-assisted instruction in statistics: A meta-analysis. *Review of Educational Research*, 81(1), 97–128. <https://doi.org/10.3102/0034654310378174>
- Striwe, M. (2016). An architecture for modular grading and feedback generation for complex exercises. *Science of Computer Programming*, 129, 35–47. <https://doi.org/10.1016/j.scico.2016.02.009>
- Striwe, M., Balz, M., & Goedicke, M. (2009). A Flexible and Modular Software Architecture for Computer Aided Assessments and Automated Marking. *Proceedings of the First International Conference on Computer Supported Education (CSEDU)*, 2, 54–61. INSTICC, 23–26 March 2009, Lisboa, Portugal.
- Striwe, M., Zurmaar, B., & Goedicke, M. (2015). Evolution of the E-Assessment Framework JACK. *Gemeinsamer Tagungsband der Workshops der Tagung Software Engineering*, 118–120.
- Wolff, A., Zdrahal, Z., Nikolov, A., & Pantucek, M. (2013). Improving retention: Predicting at-risk students by analysing clicking behaviour in a virtual learning environment. *Proceedings of the Third International Conference on Learning Analytics and Knowledge, LAK '13*, 145–149, ACM. <https://doi.org/10.1145/2460296.2460324>
- Xiong, X. Pardos, Z. A., & Heffernan, N. T. (2011). An Analysis of Response Time Data for Improving Student Performance Prediction, URL <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.374.1577>