Understanding Dropout & Repetition in Rwanda

USING PREDICTIVE SCORECARDS TO IDENTIFY AT RISK STUDENTS

October, 2017

laterite
DATA | RESEARCH | ADVISORY

Contents

Using Predictive Scorecards to Identify At-Risk Students

Acknowledgements

The analysis contained in this brief builds upon the original paper authored by Ben Brockman and Vincent Vanderputten, of the John F. Kennedy School of Government, Harvard University, who, in early 2017, developed a predictive scorecard for repetition based on data from Rwanda's fourth Integrated Household and Living Conditions Survey (EICV4).[1] The current analysis herein focuses on developing predictive scorecards using data collected as part of the "Understanding Dropout and Repetition in Basic Education in Rwanda" study, which is funded by UNICEF. This analysis was conducted by Vincent Vanderputten in partnership with Laterite Ltd.

We thank UNICEF for supporting this effort.

[1] Ben Brockman and Vincent Vanderputten. "Prediction and Prevention: Reducing Repetition and Dropout of Primary Students in Rwanda". *Second Year Policy Analysis, John F. Kennedy School of Government, Harvard University*. 2017.

Using Predictive Scorecards to Identify At-Risk Students

1.0. Executive Summary

Identifying students at an increased risk of grade-repetition or dropout is an essential first step in deploying targeted pre-emptive interventions. In the Rwandan context, repetition and dropout are educational events that are not only driven by the educational performance of children, but also by external factors related to the socio-economic situation of children's households. School grades or absences alone, are not sufficient to identify children at a high risk of dropout or repetition. To effectively pre-empt dropout and repetition, schools need to have a mechanism to identify children at-risk.

In this brief, we show how simple predictive scorecards can be used to identify children at-risk prior to the actual events of dropout and/or repetition taking place. By predictive scorecard, we refer to a simple-to-complete and short questionnaire with high predictive power. The predictive power of the scorecard is obtained using machine learning methods (called classifiers), which make it possible to identify the set of questions which do the best job in predicting whether children will repeat or dropout in the future. The advantage of such a scorecard over ad-hoc methods at the school level – such as identifying students at risk by looking at average grades or the number of absences – is that machine learning provides testable results, that make use of all available data.

Predictive scorecards could provide a standardized way for schools to use data to better target resources to the students who are most at risk in a simple, efficient and transparent manner. Predictive scorecards can provide a way for schools and teachers to categorize the risk level for individual children in a classroom. Schools, in collaboration with the community, can then use these scorecards to better target the use of human, financial and material resources to prevent dropout and grade repetition from happening. The use of a mechanism to identify children at risk early-on enables a more student-centric approach to tackling the problems of repetition and dropout.

This policy brief introduces a proof-of-concept predictive scorecard for grade repetition. It focuses less on the technical aspect of running the required machine learning algorithms, and more on the use and purpose of such a scorecard. It discusses how such a scorecard could be used to target preventative measures against grade repetition and shows how predictive scorecards are created. It also discusses how a predictive scorecard can only be successful in the context of a consultative process with multiple stakeholders and in combination with targeted pre-emptive initiatives. Equally important to the predictive power of a scorecard, is the question of how accessible it is and how feasible it is to collect the required information. Ensuring the scorecard is grounded in the local context is key.

We recommend that a predictive-scorecard approach - in conjunction with specially designed interventions to help schools and communities reduce dropout and repetition rates - be piloted by Rwanda's Ministry of Education, with external funding, as part of a randomized-controlled-trial.

1.1. Introduction to Predictive Scorecards

A predictive scorecard is a simple-to-use measurement tool created for the purpose of identifying and targeting a particular group. The Progress Out of Poverty Index is a well-known example of this. The PPI is a set of scorecards that are used to calculate the likelihood that a specific household is living below the poverty line in a given country. For example, the Rwanda Progress Out of Poverty Index is developed using Rwanda's integrated household living conditions survey (EICV) and constitutes a scorecard of 10 simple questions. Financial institutions across the globe use predictive scorecards to calculate the likelihood that a loan applicant would default; targeting their financial services to those that are predicted to not be at risk.

These scorecards are created using data analysis techniques based on machine learning algorithms. In this case, the machine learning algorithm "learns" which questions do the best job in predicting whether children will repeat or dropout in the future. The result of this analysis is a set of questions with categorical answers where each answer is assigned an integer score. The sum of the scores for all questions on a scorecard gives the total score for a given individual. This provides a relative measure of how likely a certain outcome is for that individual. For example, a repetition scorecard will give a measure of the risk that an individual child will repeat their current grade relative to the population of all enrolled children.

This study focuses on identifying the factors that do the best job in identifying which individual students will repeat or dropout in the future. It is important to note that the predictive analysis conducted to create a scorecard has no causal element. If including a specific factor in the model provides a more accurate prediction of whether a child will repeat, it does not follow that this is what causes the child to repeat. For example, the repetition scorecard identifies that whether the household owns 2 or more smartphones can be used to predict repetition. This does not mean grade repetition is caused by the lack of smartphones in the household. What this could mean is that smartphone ownership is correlated with other factors that cause grade repetition (household income, employment status, parental education, etc). We do not know - and for the purposes of prediction do not need to know - what these casual factors are.

In this study we employed predictive analytic techniques using primary data collected from the *"Understanding Dropout and Repetition in Basic Education in Rwanda"* study, funded by UNICEF. This study involved four sets of survey instruments, including child, household, community and school-level surveys. Child-level data was connected to household, community and school data through unique identifiers; combined these datasets included several thousand potential predictors. Using machine learning algorithms, we identified the set of predictors with the highest predictive power. The result is a 15-question scorecard for predicting repetition.

In this study we introduce a proof-of-concept predictive scorecard that aims to show how such a tool could be used to identify children with the highest risk of future repetition. We focus less on the actual implementation of the learning algorithm, and more on the purpose and potential implementation of such a scorecard.

1.2. Repetition: Predictive Scorecard

In this sub-section we introduce an example of the predictive scorecard we obtain after running a machine learning algorithm on data from the *"Understanding Dropout and Repetition in Basic Education in Rwanda"* study.

Using this scorecard, and 15 simple questions, teachers can identify students who are at an elevated risk of repeating the grade they are currently enrolled in. Figure 1.1 shows our one-page proof-of-concept predictive scorecard. In practice, the teacher would fill in this card for each child in their class. Some of the information can be filled in without follow up while some needs to be answered after discussion with the student and/or parents. Each response corresponds to either 0 points or a specific number of points towards a final score. The higher the score, the higher the repetition risk for the individual is.

Figure 1.1. Proof-of-concept Scorecard for Repetition

Draft Scorecard: Student Repetition Risk Student Name: ______________________

Q1) Circle the cell on the "Score" row, based on pupil's grade, and write down the score. **Scores:**

Grade ➔	**P1**	**P2**	**P3**	**P4**	**P5**	**P6**		
Score	5	0	0	4	12	15	➔	______

Q2) For each of the 4 grade-appropriate tests, circle the score to add in the right grade column.

grade out of 10 ➔	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**		
1) Addition Test	8	7	6	6	5	4	3	2	2	1	0	➔	______
2) Subtraction Test	2	2	2	1	1	1	1	1	0	0	0	➔	______
3) Multiply Test	12	11	10	8	7	6	5	4	3	1	0	➔	______
4) Literacy Test	13	12	10	9	8	7	5	4	3	1	0	➔	______

Q3) Does the pupil live in an urban or rural area? (*In this and next questions, circle answer and copy to score column.*)	Urban Rural	6 0	______
Q4) Can the pupil ask for homework help from at least one person?	Yes No	0 2	______
Q5) Are there other schools offering primary schooling in the community?	Yes No	0 3	______
Q6) Did the pupil ever repeat grade 1 (P1)?	Yes No	0 2	______
Q7) Has the pupil repeated his/her current grade?	Yes No	0 12	______
Q8) How may computers does the pupil's family have?	Zero One or more	10 0	______
Q9) Does the household own 2 or more smartphones?	Yes No	0 6	______
Q10) At home, who helps the pupil ***the most*** with homework?	Siblings Parents	0 2	______
Q11) Do the parents force the pupil to go to school even when he/she does not want to?	Yes No	3 0	______
Q12) Are the parents satisfied with pupil's performance at school?	Yes No	0 3	______
Q13) When the pupil experienced a problem at school, did his/her teachers talk to the parents about it?	Yes No	0 3	______
Q14) When the pupil missed school, did someone from your school contact the pupil or family?	Yes No	0 1	______
Q15) Is either Fishing or Factory Work one of the three main sources of employment in the community?	Yes No	8 0	______

TOTAL RISK SCORE (Sum of all values in the score column)**:** ______

As we would expect grade and test scores are powerful predictors of repetition risk. An interesting variable that comes out as a predictor for repetition risk is the presence/absence of other primary schools in the community. The presence of other primary schools in the community is associated with a lower risk of repetition.

The total score that an individual student has cannot be interpreted as a direct probability of repeating. What is more important for identifying at-risk children is the relative position of that child's score. By relative position we mean how the score relates to the average score for a child in Rwanda or for children in this child's specific district, school or classroom. Students with higher scores are more likely to repeat than students with lower scores.

We use the repetition scorecard from Figure 1.1. to calculate a risk score for each of the students in the dataset. The distribution of predicted repetition risk scores is shown in Figure 1.2. We find a range of scores from 28 to 88 out of the maximum of 111. We also show whether the students with these scores actually repeated or dropped out. Hardly any students with scores below 50 ultimately ended up repeating. Most repeaters have scores above 60, even though in practice there are also many students with scores above 60 who ultimately did not repeat.

Figure 1.2. Distribution of repetition risk scores disaggregated by ultimate outcome (whether the child repeated or not)

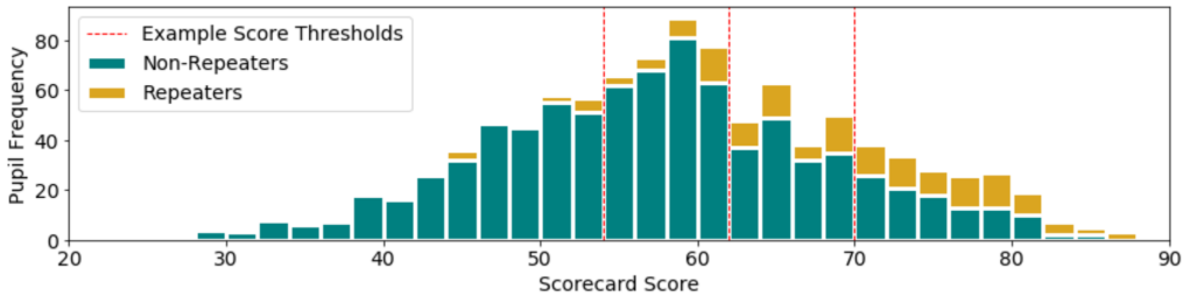

This leads to the important question of establishing the appropriate threshold to identify children with the highest risk of repetition. After what threshold do we consider a child to be at a high risk of repetition?

1.3. Repetition: "At-Risk" Students

Teachers need to know which students are "at-risk" of repetition to be able to take preventative action: setting a threshold risk score for "at-risk" children provides a simple way of determining when to take preventative action. Setting too high a threshold will mean that many children at risk of repetition will be missed (under-coverage). In a high-threshold scenario, resources will not be wasted on targeting children not at risk of repetition. Setting too low a threshold will mean that almost all of the children at risk of repetition will be targeted, but many children not at risk of repetition will also be targeted (leakage). This is inefficient from an resource perspective. This trade-off between under-coverage and leakage is shown in the bottom panel of Figure 1.3.

Figure 1.3. Top: Distribution of repetition risk scores disaggregated by ultimate outcome (whether the child repeated or not). Bottom: Prevalence of under-coverage and leakage with different "at-risk" thresholds

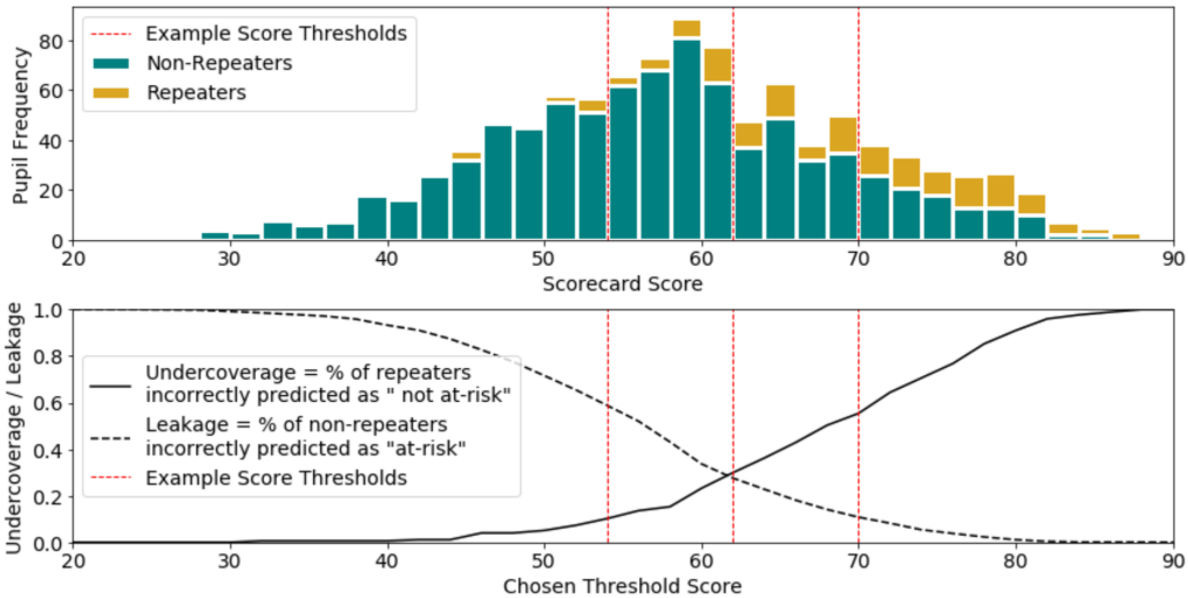

Deciding on a threshold for assigning students as "at-risk" or "not-at-risk" is not an automated process but a decision taken by testing different thresholds. Setting a lower threshold will reduce under-coverage (failing to identify students at risk of repetition) but at the same time will increase leakage (incorrectly suggesting students are "at-risk" when they are not). We test three example thresholds for our repetition scorecard and calculate the prevalence of under-coverage and leakage for a hypothetical classroom of 100 students. We find that a threshold of 62 points for our repetition scorecard leads to under-coverage of 29.9% and leakage of 27.6% (Figure 1.4.). This means that approximately 30% of children who ultimately repeat failed to be identified as "at-risk" by our scorecard while approximately 30% of children who did not ultimately repeat were wrongly identified as "at-risk".

Figure 1.4. Performance measures of three example thresholds in a hypothetical representative classroom of 100 students of which 17 students ultimately repeated

*The impossible **perfect** model*	Predicted as at-risk	Predicted as **not** at-risk
Actual repeater	17	0
Actual **non**-repeater	0	83

*Scorecard with threshold of **54+***	Predicted as at-risk	Predicted as **not** at-risk
Actual repeater	15	2
Actual **non**-repeater	50	34

*Scorecard with threshold of **62+***	Predicted as at-risk	Predicted as **not** at-risk
Actual repeater	12	5
Actual **non**-repeater	23	60

*Scorecard with threshold of **70+***	Predicted as at-risk	Predicted as **not** at-risk
Actual repeater	8	9
Actual **non**-repeater	9	74

Indicators to ****both**** minimize ➔	Leakage	Under-coverage
*The impossible **perfect** model*	0%	0%
*Scorecard with threshold of **54+***	58.8%	10.2%
*Scorecard with threshold of **62+***	27.6%	29.9%
*Scorecard with threshold of **70+***	10.8%	55.4%

In practice, the threshold needs to balance this trade-off and take into consideration the resources and interventions available. If the intervention is low-cost, affordable and has no negative consequences it may be better to set the threshold lower. If the intervention carries a high cost then it may be better to set a higher threshold so that expensive resources are not wasted on students who were not ultimately going to repeat even without the intervention.

One idea for the use of scorecards for targeting preventative measures against dropout and repetition, would be to have two thresholds. This would classify children into three categories, “not-at-risk”, “medium-risk” and “high-risk”. In this way preventative measures where leakage would be less consequential such as supplemental instruction or homework help groups could be put in place for the medium risk group while measures including financial support or fee exemptions are reserved for the “high-risk” group only.

1.4. Repetition: Using the Scorecard

The repetition scorecard presented above (Figure 1.1) is a proof-of-concept scorecard that shows how predictive analysis can be used to identify students at risk of repetition. In practice, the process of creating a final implementation-ready scorecard for dropout or repetition is an iterative, participatory and collaborative effort between multiple stakeholders including central government officials, local community officials, head-teachers and teachers.

The next steps would be to assess how easy it would be in-practice to collect the necessary information for the scorecard. Below we detail some aspects of implementation that must be considered:

i) **Who would be responsible for collecting the scorecard data**? The two most likely candidates are teachers or community leaders. Logistically teachers may be the best option. This would also improve teacher buy-in by providing a sense of ownership but comes with a risk that scores may be influenced by the teachers' pre-conceived ideas about their pupils or mistrust of the scorecard model.

ii) **Who would data need to be collected from?** Would information from the students themselves suffice or would parents need to be consulted?

iii) **Are the predictor variables appropriate?** If necessary at this point we would exclude variables that might be deemed too difficult to collect and re-run the predictive analysis as many times as necessary with the end product being a scorecard with as high a predictive accuracy as possible but also practically feasible for the Rwandan context.

iv) **Is one national scorecard is sufficient?** What benefit would there be to creating provincial or district level scorecards? Should there be separate scorecards for urban and rural areas?

v) **How often would scorecards be filled out?** Would this happen once at the beginning of the school year or once per school term?

vi) **How would the information be stored? What would be the requirements in schools for making sure scores are tracked over time for the same pupils? Even when those pupils move schools?** Electronic scorecard data collection paired with the "One Digital Identity Per Child" goal proposed under the ICT in Education Policy (MINEDUC 2016) may be the most practical way of doing this.

Some of these questions would be best answered through stakeholder discussions but some would require piloting the scorecard in a small area to begin with and adjusting the scorecard protocol accordingly. Another important thing to consider would be whether we expect any unintended negative effects for children who are categorized as "at-risk" for example, if this would cause stigma in the classroom or in the community.

It is important to remember that having an accurate, simple-to-use scorecard for identifying which students are at-risk of repeating or dropping out from school is ultimately only useful if there are preventative measures in place for the schools to implement with the targeted children.

There is no template for what should come next after at-risk students are identified, however, some example preventative measures that might be implemented by schools are:

- **Teachers to use their individual discretion to support "at-risk" students on a case-by-case basis**.

- **Create a standardised program of support that is put in place for all at-risk students.** Such a program might include for example, supplemental instruction, homework help groups and financial support such as parental exemption from school feeding fees. As discussed in the assigning thresholds section of this brief, different levels of intervention can have different risk thresholds with financial support being reserved for only the highest risk category.

- **Implement peer-to-peer mentoring** whereby students with a very low risk of repetition are paired with at-risk children for classroom or homework activities.

- **Introduce specific performance targets for "at-risk" children**: children at-risk of repetition can be identified at the start of the year and teachers can be set performance targets for ensuring these students are ready to progress at the end of the year (measured by end of year test scores). This could be paired with teacher incentives.

As discussed previously in this brief the predictor variables for a scorecard provide no causal information on repetition or dropout. This is a very important concept that would need to be understood by all education agents who put such a scorecard into practice.

This is necessary to avoid misconceptions that lowering a score based on changing the child's answers to some of the predictor variables would not change the actual likelihood that the child is at-risk of the negative outcome (repetition or dropout). To illustrate this with an example the scorecard in Figure 1.1. includes the question "Do the parents force the student to go to school even when he/she doesn't want to?" with points added to the repetition risk for students where the answer is "yes". Instructing parents to not force their children to go to school if they don't want to might lead to a small reduction in the average class scorecard score but it would not make those children less likely to actually repeat.

By extension, an overall reduction or increase in the occurrence of the predictive factor within the population will not directly impact actual probabilities of dropout and repetition but will make the scorecard less accurate at predicting. This is because the population that the model is trying to predict the behavior of has changed from the one that was used to create it. For this reason, the model and set of questions that make up a scorecard need to be updated regularly to maintain and improve predictive power over time.

1.5. Dropout: Requirements for a Predictive Scorecard

We investigated whether it is possible to develop a Predictive Scorecard for dropout and concluded that data limitations make it impossible to create a dropout scorecard at this stage.

We started by looking at what variables are best at predicting dropout. These are detailed in Table 1.1. This shows that the two most important predictors for whether a child is at risk of dropping out are age and grade. Firstly, as age increases risk of dropout also increases. As well as this, simply being enrolled in Primary 6 is by far the strongest predictor of dropout.

Table 1.1. Predictor variables and coefficients selected

Predictor Variable	Description	Coefficient
age	Age of the child	0.7554
repeat_grade4	The child has previously repeated Primary 4	0.0027
repeat_grade5	The child has previously repeated Primary 5	-0.0225
chair_number	Number of chairs that the child's household owns	-0.0711
blankets_number	Number of blankets that the child's household owns	-0.0135
serv_avail_1	The community has electricity	0.1074
drop_prevyr	The child has re-entered the school system this academic year having been out of school previously	-0.0638
m4_parents_16_often	If the child is doing well teachers would tell the parents/guardians often	-0.0188
inv_teacher_very_often	The parents/guardians met with the teacher very often to discuss their children's education	0.0182
send_school_disagree	The parent disagrees that sending children to school is a waste of time	0.0003
value_school_money_agree	Parent agrees that the main value of sending children to school is to allow them to improve their future income	0.0071
asp_max_want_dontknow	The parent doesn't know what level of education they would like their children to achieve	0.0314
grade_5	The child currently enrolled in Primary 5	0.0119
grade_6	The child currently enrolled in Primary 6	0.4512
sm_construction	Construction work is one of the three main sources of employment in the community	0.2094

Model estimated using LASSO regularization on normalized data and specifying a model with 11 variables. (AUC = 0.9342)

Interestingly, children who have repeated Primary 5 are less likely to be at risk from dropout, as are those who have dropped out and re-entered in the immediate past (one year previously). Children who come from communities where there is electricity or where construction is a main source of employment are more at risk of dropping out. One important thing to remember with these predictor variables is that their predictive importance does not imply any underlying causation. The presence of electricity in a community is not causing children to dropout from school but we find that communities with electricity happen to be at higher risk of student dropout.

The main issue with trying to create a scorecard to predict a students' risk of dropping out is that the data we are using has been collected after-the-fact. This is very limiting for predictive analysis as we know that after dropping out a child's situation is very different than it was before. For example, it has been shown in the main analysis for this study that after dropping out from school children tend to start working, work longer or spend many more hours engaged in household chores than their enrolled counterparts. This means that our socio-economic data on children that have dropped out during the 2016 school year is not likely to give us an accurate model for predicting which enrolled children are likely to dropout in the future. While this model is very accurate at predicting who will dropout that is because it is using the data of children who have already dropped out.

Another issue with creating a scorecard for school dropout is that dropout is a rare event meaning that even with a large dataset the number of dropouts captured is relatively low and we are at risk of creating a model that is sample-specific and not generalizable. Additionally, when the main predictor variable is whether the child is enrolled in Primary 6, it is extremely difficult to make a scorecard that is nuanced enough to capture "at-risk" children in other grades or distinguish "at-risk" and "not-at-risk" students within Primary 6 enrolled children. With dropout itself being a rare event and dropout from Primary 1 to Primary 5 even rarer it was not possible to create such a scorecard with the dataset from this study.

To create a dropout scorecard would require a panel dataset, with data on children's individual, school and ideally also household situation collected on a regularly basis, yearly, for example, with the same individuals each year. Such a dataset would overcome the issue of not knowing an out-of-school child's situation prior to the dropout event and would, over time capture enough dropout events to better generalize. An expanded and regularly updated EMIS including this type of information would form the ideal basis for creating predictive models for dropout in the future.

1.6. Developing the Predictive Scorecards

A successful predictive model for dropout or repetition needs to be accurate, correctly identifying as many at-risk students as possible, while remaining simple and limiting the data collection effort required for their regular use.

1.6.1. Data Requirements

Creation of a predictive scorecard requires data that is statistically representative of the population of interest, in our case children enrolled in primary school at the beginning of the 2016 school year. This initial input dataset should include data at the individual level covering as many socio-economic areas as possible, (demographics, health, education, housing, services, income and assets, etc.), as well as the ultimate outcome of the 2016 school year for each child i.e. whether they were promoted, repeated or dropped out at some point before the beginning of the 2017 school year.

After the initial creation of a scorecard the predictive analysis carried out here needs to be re-run with each new year of data to keep up with changes in the population and keep improving the predictive accuracy of the model.

1.6.2. Model Training and Validation

The dataset is first split into two, one dataset is set aside for a later testing phase while the other is used for the initial training and validation of the model. Training and validation is the process whereby predictor variables - the variables that will end up in the scorecard - are selected. In order to select the variables that combine to provide the best predictive model for a binary (yes/no) outcome such as dropout or repetition we use a quantitative technique called logistic regression with LASSO regularization. Logistic regression is a statistical technique which can classify a binary variable—that is, where it can take only two values, "0" and "1", which represent outcomes such as pass/fail, win/lose, repeat/no-repeat, dropout/no-dropout. The logistic model is used to estimate the probability of the outcome being "1" based on predictor variables. The LASSO regularization step is used to avoid overfitting (including too many predictor variables).

The training and validation step is taken for different levels of regularization leading to models with different numbers of predictor variables. The purpose of this exercise is to explore the trade-off between simplicity and predictive power. A complex model with many predictor variables might result in higher predictive accuracy but also increases the cost and accessibility of the final scorecard. We do not want to create a scorecard with maximum accuracy if it is ultimately too complex for the agents who we intend it to be used by, in this case primary school teachers. Similarly, a model with only a few predictor variables will make data collection and the scorecard use easier and cheaper but may result in fewer accurate predictions.

We compare different models using a statistic called the AUC (Area Under the Curve), which is a measure of this trade-off between limiting complexity and maximizing accurate predictions. The AUC is the probability that the model will correctly identify the dropout/repeater from a randomly drawn pair in which one child dropped out/repeated and one child did not. A perfect prediction model has an AUC of 1 meaning that all at-risk students are correctly predicted as "at-risk" and none of the not-at-risk students are predicted as being "at-risk". An AUC of 0.5, describes a model that randomly guesses whether a student is at-risk or not. We ultimately want to create a model that has an AUC as close to 1 as possible.

1.6.3. Creating the Scorecard

The final scorecard is created by processing the raw coefficients of each of the predictor variables. Adding the coefficients weighted by the value of a pupil's predictor variables gives a number that is not a probability, but maps to it 1-to-1 after applying the logit function. Hence the final score a model gives for a student is a level of risk, but not an absolute probability of dropout/repetition. Coefficients are processed by rounding to the nearest integer, then assigning them to the correct category of the predictor variable (yes or no) and lastly, formulating the correct wording of the question for inclusion in the scorecard.

1.6.4. Assigning the "at-risk/not-at-risk" Threshold Score

Once the scorecard has been finalized, the last step is to choose a cut-off score above which a student would be flagged as "at-risk". This can be done in a number of ways. For example, we could take the 10% of students with the highest scores as "at-risk". A teacher with a class of 50 students may focus extra attention on the 5 students with the highest risk scores. In contrast, a specific score threshold might be set at the central level and schools be instructed such that all children with a score above X are targeted with preventative measures programs.

1.6.5. Testing Scorecard Performance

After finalizing the scorecard, the second dataset, created and set aside during the initial split (test data), is used to test the predictive performance of the scorecard. Given a score for each student and a chosen threshold, we compare which students were flagged as at-risk or not based on their scores with whether those children actually ended up being dropouts/repeaters by the end of the 2016 school year. During the testing phase each individual in the dataset is categorized into one of the following four categories (examples are given in reference to the repetition scorecard):

- **True positive:** The child was predicted to be "at-risk" by the scorecard result and repeated their 2016 grade
- **False positive:** The child was predicted to be "at-risk" by the scorecard result but did not repeat their 2016 grade
- **True negative:** The child was predicted to be not-at-risk by the scorecard result and did not repeat their 2016 grade
- **False negative:** The child was predicted to be not-at-risk by the scorecard result but repeated their 2016 grade

We look at two diagnostic measures relating to these categories; under-coverage and leakage.

- **Under-coverage** refers to the prevalence of false negatives, the likelihood that we will fail to identify at-risk students and therefore not target them for preventive measures.
- **Leakage** refers to the prevalence of false positives, the likelihood that we will incorrectly identify not-at-risk students as at-risk leading to wasted resources.