

1 How To Use This Document

Highly regulated industries, such as banking and insurance, must comply with government regulations for model validation before a model can be put into production. This includes creating robust model development documentation. DataRobot automates the generation of model documentation, expediting the process required for regulatory compliance and following best practice for reducing model risk.

This document is split into two components: those sections that are automatically produced by DataRobot and those that require further input by the user. The sections in *blue italicized font* include specific instructions for the documenter and require additional user input of organization-specific information, such as business use cases, data sources, and implementation details. Once the sections are complete, remove the instructions. The remaining sections in non-blue italicized font are automatically populated by DataRobot and require no further input.

Copyright ©2023, DataRobot, Inc.

The logo for DataRobot, featuring the word "Data" in a bold, black, sans-serif font, followed by "Robot" in a bold, blue, sans-serif font.

Table of Contents

- 1 How To Use This Document
- 2 Model Performance Overview
- 3 Model Features Summary
 - 3.1 Model Features and Summary Statistics
 - 3.2 Data Quality Handling Report
- 4 Model Summary and Description
- 5 Feature Effects
- 6 Feature Impact Chart
- 7 Feature Impact Table
- 8 Validation Testing and Stability
 - 8.1 Cross Validation Scores
- 9 Model Results
- 10 Bias and Fairness

2 Model Performance Overview

As an additional layer of model validity, DataRobot not only evaluated the statistical metrics underlying the model, but also performed testing on in-sample records.

The performance metric used for this project was LogLoss. The model performance results are presented below for in-sample testing:

| Scoring Type | Score (LogLoss) |
|------------------|-----------------|
| cross_validation | 0.4685* |
| holdout | 0.0679* |
| validation | 0.4568* |

3 Model Features Summary

Below are two tables. The first contains a list of the final set of model feature variables, as well as summary statistics for the Regularized Logistic Regression (L2) model. The second table contains a detailed analysis of missing values.

The Model Features and Summary Statistics table provides a brief overview of the summary statistics of model features. This includes Feature Name, variable type (Var Type), number of unique values (Unique), Number of missing values (Missing), Mean, Standard Deviation (Std Dev), Median, Minimum Value (Min), Maximum Value (Max) and Assessment of target leakage risk (Target Leakage).

3.1 Model Features and Summary Statistics

| Feature Name | Var Type | Unique | Missing | Mean | Std Dev | Median | Min | Max | Target Leakage |
|-----------------------|-------------|--------|---------|----------|----------|----------|--------|-----------|----------------|
| water_tech_clean | Categorical | 15 | 734 | N/A | N/A | N/A | N/A | N/A | Low |
| water_source_clean | Categorical | 6 | 519 | N/A | N/A | N/A | N/A | N/A | Low |
| management_clean | Categorical | 8 | 467 | N/A | N/A | N/A | N/A | N/A | Low |
| pay_clean | Categorical | 6 | 10017 | N/A | N/A | N/A | N/A | N/A | Low |
| age_in_years | Numeric | 4961 | 1 | 9.92 | 9.98 | 7.74 | -26.22 | 67.81 | Low |
| distance_to_primary | Numeric | 20761 | 0 | 21192.78 | 16070.78 | 18546.53 | 0.0065 | 83299.34 | Low |
| distance_to_secondary | Numeric | 20761 | 0 | 13208.12 | 14802.36 | 8602.75 | 0.23 | 98940.72 | Low |
| distance_to_tertiary | Numeric | 20761 | 0 | 5980.29 | 6956.89 | 3683.703 | 0.075 | 50760.63 | Low |
| distance_to_city | Numeric | 20761 | 0 | 69377.43 | 35751.55 | 64787.93 | 373.15 | 163868.54 | Low |

| | | | | | | | | | |
|---------------------------------------|---------|-------|-----|------------|------------|-----------|------------|------------|-----|
| distance_to_town | Numeric | 20761 | 0 | 10529.058 | 7682.69 | 9101.14 | 13.24 | 51911.46 | Low |
| bgs_stor | Numeric | 4 | 0 | 2504.92 | 4313.76 | 0.0 | -1.0 | 10000.0 | Low |
| bgs_recharge | Numeric | 213 | 0 | 93.59 | 17.78 | 92.72 | 0.0 | 132.59 | Low |
| water_risk | Numeric | 56 | 0 | 2.28 | 0.58 | 2.14 | -1.0 | 3.84 | Low |
| rwi | Numeric | 457 | 0 | -0.31 | 0.39 | -0.39 | -1.34 | 1.33 | Low |
| assigned_population | Numeric | 2719 | 477 | 574.0603 | 1211.508 | 237.0 | 0.0 | 33044.0 | Low |
| local_population | Numeric | 4753 | 477 | 1740.56 | 3651.74 | 687.0 | 0.0 | 78642.0 | Low |
| pressure | Numeric | 2986 | 968 | 2.0 | 4.26 | 0.82 | 0.0033 | 123.62 | Low |
| crucialness | Numeric | 14614 | 968 | 0.45 | 0.3 | 0.35 | 0.016 | 1.0 | Low |
| population_1km | Numeric | 11241 | 0 | 1617.58 | 3300.88 | 661.86 | -0.0 | 72264.23 | Low |
| men_ratio_1km | Numeric | 2064 | 0 | 0.502 | 1.63 | 0.49 | -22.5 | 146.5 | Low |
| women_of_reproductive_age_ratio_1km | Numeric | 2066 | 0 | 0.24 | 0.703 | 0.24 | -19.5 | 71.0 | Low |
| youth_ratio_1km | Numeric | 2036 | 0 | 0.19 | 0.91 | 0.2 | -91.75 | 27.5 | Low |
| elderly_ratio_1km | Numeric | 2069 | 0 | 0.07 | 0.1004 | 0.062 | -4.062 | 9.34 | Low |
| children_under_five_ratio_1km | Numeric | 2003 | 0 | 0.14 | 0.43 | 0.14 | -39.0 | 26.0 | Low |
| population_10km | Numeric | 11874 | 0 | 46712.33 | 150721.18 | 26338.94 | 39.84 | 2923419.5 | Low |
| men_ratio_10km | Numeric | 9268 | 0 | 0.49 | 0.013 | 0.49 | 0.45 | 0.53 | Low |
| women_of_reproductive_age_ratio_10km | Numeric | 9281 | 0 | 0.24 | 0.0091 | 0.24 | 0.22 | 0.27 | Low |
| youth_ratio_10km | Numeric | 8685 | 0 | 0.2 | 0.0059 | 0.2 | 0.17 | 0.207 | Low |
| elderly_ratio_10km | Numeric | 9068 | 0 | 0.07 | 0.013 | 0.062 | 0.046 | 0.092 | Low |
| children_under_five_ratio_10km | Numeric | 8909 | 0 | 0.15 | 0.012 | 0.14 | 0.13 | 0.2 | Low |
| population_100km | Numeric | 1413 | 0 | 2746364.82 | 1774640.32 | 2217927.5 | 382840.406 | 10121682.0 | Low |
| men_ratio_100km | Numeric | 1412 | 0 | 0.49 | 0.0085 | 0.49 | 0.48 | 0.52 | Low |
| women_of_reproductive_age_ratio_100km | Numeric | 1412 | 0 | 0.24 | 0.011 | 0.24 | 0.22 | 0.28 | Low |
| youth_ratio_100km | Numeric | 1413 | 0 | 0.19 | 0.0087 | 0.19 | 0.17 | 0.206 | Low |
| years_since_report | Numeric | 4920 | 0 | 9.76 | 10.0 | 7.33 | -26.22 | 67.81 | Low |

The last column in this table is an assessment of target leakage risk. DataRobot automatically tests for target leakage on a per-feature basis during the Autopilot process. Target leakage, sometimes called data leakage, occurs when a model is trained using a dataset that includes information that would not be available at the time of prediction. This can produce overly optimistic model performance results during training, given a feature will near-completely describe the target (e.g., the number of late payments on a loan as a predictor for loan default at loan application date.)

DataRobot tests for target leakage risk using Alternating Conditional Expectation (ACE) to measure the association between each feature and the target; the ACE score is normalized using the project optimization metric so that its value is in the range [0,1]. If above a certain threshold (see below), DataRobot will create a new feature list with those features flagged and possibly removed, and the

user is notified by a banner in the user interface during modeling. Notably, because the definition of target leakage is directly tied with prediction time and not strength of association between a feature and the target, it's possible for DataRobot to not identify all sources of target leakage. Therefore, to reduce the risk for potential target leakage in the feature list, it's important to apply subject matter expertise.

The thresholds for target leakage risk are based on a normalized ACE score:

- High risk: > 0.975, flagged and removed
- Moderate risk: > 0.85, flagged but not removed
- Low risk: < 0.85, no action

The following table provides a summary of missing values. It includes the name of the feature, its type, a summary of the missing value count (both number of rows and as a percentage), and information on the type of imputation applied to the feature.

3.2 Data Quality Handling Report


| Feature Name | Var Type | Missing Count | Missing Percentage | Imputation Name | Imputation Description |
|-----------------------|-------------|---------------|--------------------|------------------|--------------------------------------|
| pay_clean | Categorical | 13780 | 54 | One-Hot Encoding | Missing indicator treated as feature |
| water_source_clean | Categorical | 2687 | 11 | One-Hot Encoding | Missing indicator treated as feature |
| pressure | Numeric | 1783 | 7 | Constant Splines | Missing indicator treated as feature |
| crucialness | Numeric | 1783 | 7 | Constant Splines | Missing indicator treated as feature |
| water_tech_clean | Categorical | 1628 | 6 | One-Hot Encoding | Missing indicator treated as feature |
| assigned_population | Numeric | 1224 | 5 | Constant Splines | Missing indicator treated as feature |
| local_population | Numeric | 1224 | 5 | Constant Splines | Missing indicator treated as feature |
| management_clean | Categorical | 630 | 2 | One-Hot Encoding | Missing indicator treated as feature |
| age_in_years | Numeric | 5 | 0 | Constant Splines | Missing indicator treated as feature |
| distance_to_primary | Numeric | 0 | 0 | Constant Splines | Imputed value: 22663.996 |
| distance_to_secondary | Numeric | 0 | 0 | Constant Splines | Imputed value: 7612.1607 |
| distance_to_tertiary | Numeric | 0 | 0 | Constant Splines | Imputed value: 3885.7849 |
| distance_to_city | Numeric | 0 | 0 | Constant Splines | Imputed value: 60357.874 |
| distance_to_town | Numeric | 0 | 0 | Constant Splines | Imputed value: |

| | | | | | |
|---------------------------------------|---------|---|---|------------------|-----------------------------|
| | | | | | 9488.2463 |
| bgs_stor | Numeric | 0 | 0 | Constant Splines | Imputed value: 0 |
| bgs_recharge | Numeric | 0 | 0 | Constant Splines | Imputed value: 95.7951 |
| water_risk | Numeric | 0 | 0 | Constant Splines | Imputed value: 2.1782 |
| rwi | Numeric | 0 | 0 | Constant Splines | Imputed value: -0.394 |
| population_1km | Numeric | 0 | 0 | Constant Splines | Imputed value: 709.7878 |
| men_ratio_1km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.494 |
| women_of_reproductive_age_ratio_1km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.2423 |
| youth_ratio_1km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.1997 |
| elderly_ratio_1km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.0625 |
| children_under_five_ratio_1km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.1403 |
| population_10km | Numeric | 0 | 0 | Constant Splines | Imputed value: 24871.707 |
| men_ratio_10km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.4933 |
| women_of_reproductive_age_ratio_10km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.2423 |
| youth_ratio_10km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.1997 |
| elderly_ratio_10km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.0625 |
| children_under_five_ratio_10km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.1413 |
| population_100km | Numeric | 0 | 0 | Constant Splines | Imputed value: 2467539.8 |
| men_ratio_100km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.4871 |
| women_of_reproductive_age_ratio_100km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.2436 |
| youth_ratio_100km | Numeric | 0 | 0 | Constant Splines | Imputed value: 0.1996 |
| years_since_report | Numeric | 0 | 0 | Constant Splines | Imputed value: 4.7845 |

4 Model Summary and Description

The particular model referenced in this document: Regularized Logistic Regression (L2). This model was developed in a project created with vb754f77d02337f3d of DataRobot. This model is denoted within DataRobot by the Project ID: 63fecfec7ecd518559c62ccb and the Model ID: 6407a2649b3a0f863ffdf767. The project was created on 2023-03-01 04:09:16.

The model development workflow process (i.e., the model blueprint) is detailed in the figure below.




A Blueprint represents the high-level end-to-end procedure for fitting the model, including any preprocessing steps, algorithms, and post-processing. It illustrates the many steps involved in transforming input predictors and targets into a model. Each element (or, “node”) in a blueprint can represent multiple steps.

The following elements connect to create the blueprint:


- Constant Splines
- One-Hot Encoding
- Regularized Logistic Regression (L2)

5 Feature Effects

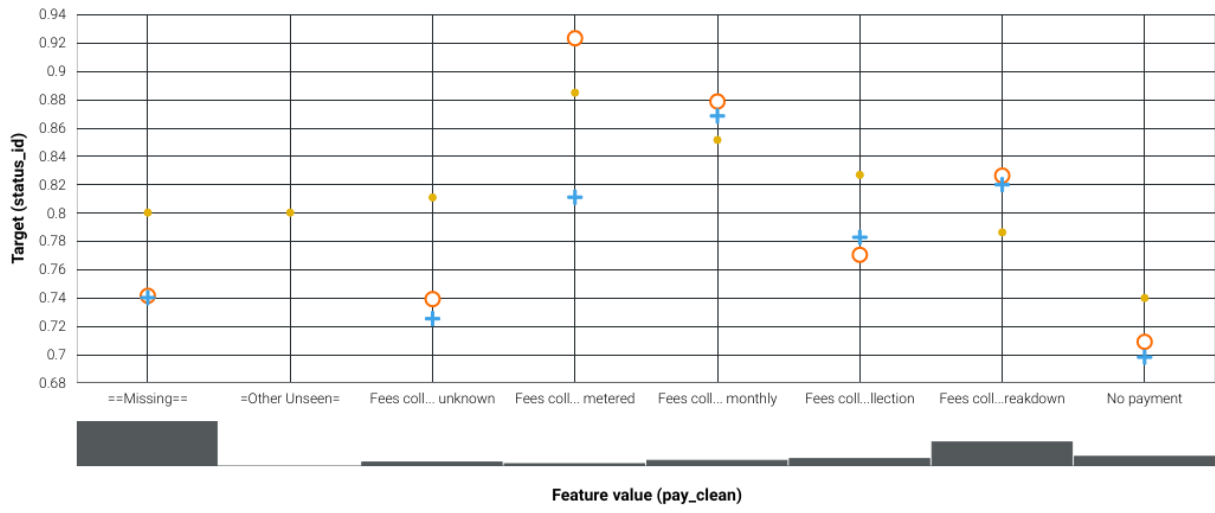
water_source_clean



age_in_years




pay_clean



● Partial Dependence + Predicted ○ Actual

6 Feature Impact Chart



7 Feature Impact Table

| Feature Name | Impact Normalized | Impact Unnormalized |
|--------------------|-------------------|---------------------|
| water_source_clean | 1.0 | 0.0492 |
| age_in_years | 0.4549 | 0.0224 |
| pay_clean | 0.2964 | 0.0146 |
| bgs_stor | 0.2502 | 0.0123 |
| elderly_ratio_1km | 0.2469 | 0.0121 |
| water_tech_clean | 0.2085 | 0.0103 |
| youth_ratio_10km | 0.1857 | 0.0091 |
| management_clean | 0.1783 | 0.0088 |

| | | |
|---------------------------------------|---------|---------|
| population_1km | 0.1646 | 0.0081 |
| local_population | 0.1388 | 0.0068 |
| women_of_reproductive_age_ratio_10km | 0.123 | 0.0061 |
| population_100km | 0.1084 | 0.0053 |
| population_10km | 0.0896 | 0.0044 |
| youth_ratio_1km | 0.0849 | 0.0042 |
| youth_ratio_100km | 0.0822 | 0.004 |
| elderly_ratio_10km | 0.0755 | 0.0037 |
| water_risk | 0.0706 | 0.0035 |
| men_ratio_1km | 0.0549 | 0.0027 |
| women_of_reproductive_age_ratio_1km | 0.0451 | 0.0022 |
| distance_to_tertiary | 0.0425 | 0.0021 |
| assigned_population | 0.04 | 0.002 |
| crucialness | 0.0399 | 0.002 |
| pressure | 0.0364 | 0.0018 |
| years_since_report | 0.0308 | 0.0015 |
| children_under_five_ratio_10km | 0.0306 | 0.0015 |
| women_of_reproductive_age_ratio_100km | 0.0262 | 0.0013 |
| bgs_recharge | 0.0257 | 0.0013 |
| men_ratio_100km | 0.0206 | 0.001 |
| rwi | 0.0105 | 0.0005 |
| men_ratio_10km | 0.0029 | 0.0001 |
| distance_to_primary | 0.0028 | 0.0001 |
| distance_to_secondary | 0.0009 | 0.0 |
| distance_to_town | -0.0041 | -0.0002 |
| distance_to_city | -0.0078 | -0.0004 |
| children_under_five_ratio_1km | -0.0136 | -0.0007 |

8 Validation Testing and Stability

To find patterns in a dataset from which it can make predictions, an algorithm must first learn from a historical example – typically from a historical dataset that contains the output variable you want to predict. However, if a model is trained too closely on its training data then it may be overfit. Overfitting is a modeling error that occurs when a model is too closely fit to training data and therefore performs poorly on out-of-sample data (data that was not used to train the model).

Overfitting generally results in an overly complex model that explains idiosyncrasies and random noise in the training data, rather than the underlying trends that the model was intended to capture. To avoid overfitting, the best practice is to evaluate model performance on out-of-sample data. If the model performs very well on in-sample data, (the training data) but poorly on out-of-sample data, that may be an indication that the model is overfit.

DataRobot uses standard modeling techniques to validate model performance and ensure that overfitting does not occur. DataRobot used a robust model k-fold cross-validation framework to test the out-of-sample stability of a model's performance. In addition to cross-validation partitioning, DataRobot uses a holdout sample to further test out-of-sample model performance and ensure the model is not overfit.

The following procedure was used during development to insure that overfitting did not occur:

- All values of the feature "fold" represent separate partitions used for cross-validation

DataRobot calculates the Cross Validation scores for each of the training data partitions or folds. The project metric used to calculate the score is LogLoss.

8.1 Cross Validation Scores

| Fold | Cross Validation Score (LogLoss) |
|--------|----------------------------------|
| Fold 1 | 0.45678 |
| Fold 2 | 0.46156 |
| Fold 3 | 0.48342 |
| Fold 4 | 0.4775 |
| Fold 5 | 0.46332 |

9 Model Results

DataRobot runs performance testing during the model development process to evaluate model results and reliability. The validation, cross-validation, and holdout (if applicable) out-of-sample performance scores are presented below, as well as the number of observations for each partition. The performance metric used for this project was LogLoss and the project included a total of 25,438 observations. An asterisk (*) next to a score, whether validation or holdout, indicates that DataRobot used in-sample predictions to derive the score. (In-samples predictions are those that include data from the validation or holdout partitions due to sample size used to build the model.)

| Scoring Type | Score (LogLoss) |
|------------------|-----------------|
| cross_validation | 0.4685* |
| holdout | 0.0679* |

10 Bias and Fairness

The Bias & Fairness feature is not a legal compliance tool. Please carefully review the product Documentation to ensure that you fully understand the functionality of the feature before using it. DataRobot recommends that you take local legal advice where you wish to rely on the results of this feature for complying with any legal obligations. Product Documentation can be found here: <https://app.datarobot.comdocs/modeling/special-workflows/b-and-f/index.html>

DataRobot's Bias and Fairness testing identifies whether the model exhibits biased behavior towards any classes in the dataset's protected features, based on the selected definition of fairness. Protected features and the fairness metric are chosen before Autopilot is started. DataRobot also provides a workflow that guides you towards an appropriate definition of fairness for the specific use case.

DataRobot's Bias and Fairness feature includes two model-level insights:

- Per-Class Bias, which shows whether the model is treating certain protected groups differently as measured by the selected fairness metric. This identifies if there is biased behavior, and if so, how that bias manifests, but not why.
- Cross-Class Data Disparity, which shows how different protected classes differ in their data distribution. This offers deeper insight into why the model is treating groups differently.

Together, these insights can help identify potential mitigation strategies for bias in the dataset and model, such as improving data collection or data sampling for specific groups.

Bias and Fairness testing was used in this project. The selected protected features were `clean_adm1`. The favorable target outcome was `No`. The selected fairness metric was `favorableAndUnfavorablePredictiveValueParity`. The fairness threshold was set at 0.8.

Favorable Predictive Value & Unfavorable Predictive Value Parity (also known as Precision and NPV Parity) measure fairness by Equal Error, and it is best suited for cases when the target is severely unbalanced. These metrics measure whether your model disproportionately discovers favorable cases or omits unfavorable cases correctly for any particular group. It is calculated relative to the model's predictions--the proportion of predicted favorable or predicted unfavorable cases that are classified correctly. Unlike True Favorable and Unfavorable Rate Parity, Favorable Predictive Value & Unfavorable Predictive Value Parity are calculated relative to the predicted class, rather than the ground truth, which allows it to capture more meaningful information when there are very few predictions for a certain class, such as in imbalanced datasets. Higher values for Favorable Predictive Value and Unfavorable Predictive Value mean that the model is more accurate for that class relative to that metric. It's important to measure both Favorable Predictive Value &

Unfavorable Predictive Value Parity together, because poor performance in either one can lead to biased outcomes, often for different stakeholders. These metrics will help you investigate how your model balances both of these types of accuracy across your protected classes.


The Per-Class Bias graph shows whether the model exhibits biased behavior across protected features. The top fairness score across each protected class is scaled to 1.0, and the fairness scores for every other class are scaled relative to that value. If the fairness score for a class crosses the selected fairness threshold, the bar for that class is shown in red. If DataRobot used in-sample predictions to derive the model's performance scores (see **Overview of Model Results**), the fairness scores were calculated using in-sample validation data.

If there is not enough data for a class, its score is still calculated, but the bar for that class is shown in gray. The heuristic for whether a class does not have enough data is the following:

- If the class has <100 rows in the validation data, then it does not have enough data.
- If the class has between 100 and 1,000 rows in the validation data, but has fewer than <10% of the rows of the majority class, then it does not have enough data.
- If the class has >1,000 rows, then it has enough data.

The following figure is the Per-Class Bias graph for each protected feature in this project:

clean_adm1



The Cross-Class Data Disparity graph shows how different protected classes differ in their data distribution, in order to understand why the model treats each class differently based on the dataset. The X-axis depicts the feature importance of each feature in the dataset, while the Y-axis shows the Population Stability Index (PSI) for that feature compared across the two selected classes of the protected feature. The higher the feature importance, the more important that feature is to the model. The higher the PSI, the more differences there are for that feature across each of the two classes.