

 Snitch suggests to **remove** current model from production

67

Model

Customer Churn Prediction

Version

1.1.2

Characteristics

- 540 Training observations
- 180 Testing observations
- 180 Production observations
- 74 Features
- Classification model
- Sequential architecture

Recommendations

Take immediate action:

- Check data drift for production observations
- Check model overfitting

Warning that requires attention:

- Check feature contribution to model's output

Summary - Model Validation Tests Performed

Feature Bias	Fail	3 features determine 95% of the output of the model
Labelling Errors	Pass	100% of the observations were correctly labelled
Data Leakage	Pass	No leakage detected in the datasets
Overfitting	Fail	Model overfit, performance on Training is 8% better than Testing
Model Simplification	Pass	Reducing complexity affects model's performance
Feature Pruning	Pass	Pruning features impacts model performance
Sensitivity to Random Noise	Pass	Model is robust to 85% random noise injection
Sensitivity to Extreme Noise	Pass	Model is robust to a 20% adversarial noise injection
Data Drift	Fail	Production observations are drifting compared to Training

Data Format

FAIL

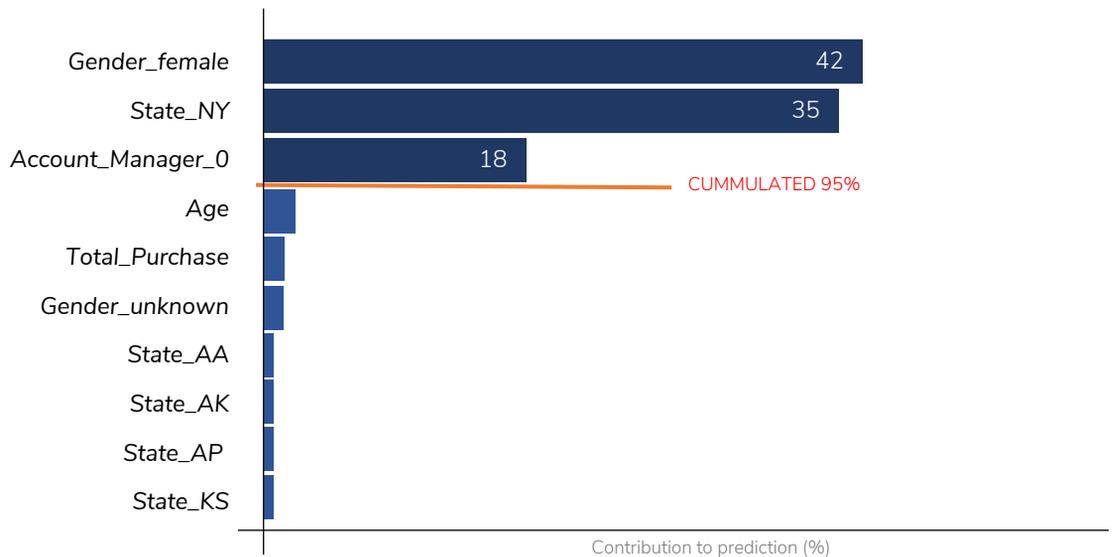
Observation

- 3 features determine 95% of the output of the model; contribution to prediction is not properly distributed across features

Recommendations

- Check if main contributors can be detailed into more granular features.
- Snitch recommends to check if predicting based on features **Gender_female, State_NY, Account_Manager_0**, can represent a liability.

Prediction Interpretation Tests Details



What is feature bias?

Feature bias occurs when a limited number of features greatly influence the output of a model. It is the hypersensitivity of a model to a specific feature. Feature bias can occur when the model is not adapted to the task it needs to perform, when data are not representative of the real world or when the data itself contain biases.

Why evaluating feature bias is important?

Unwanted feature bias can be ethically problematic: a model mainly considering the sex or the ethnicity of an individual could be perceived as sexist or racist. On the technical aspect, a model basing its predictions on simply two features is less robust: what happens when those features are not available for an individual?

How Snitch evaluates feature Bias

Feature bias can be detected by establishing the contribution of each feature when the model makes a prediction. Snitch leverages the concept of Shapley values, originally from cooperative game theory. See our [documentation](#) for more details

 **WARNING: Data is drifting**

Data Format

FAIL

Observations

- Production data contains 2 new features
- Feature **Total_Purchase** has unexpected negative values.

Recommendations

- Features **is_premium** and **is_not_premium** were not part of the provided training data. Investigate cause why these new columns have appeared.
- Investigate cause for **Total_Purchase** negative observation.

Population Stability Index

FAIL

Observation

- 2 features demonstrate major drift.

Recommendation

- Please investigate cause for drift in **Total_Purchase** and **State_NY**

Are the production observations distinct from training observation ?

FAIL

Observation

- Snitch is able to distinguish the production observations from the training observations

Recommendation

- Production observation distributions do not match training observation distributions, please check production data and/or retrain model.

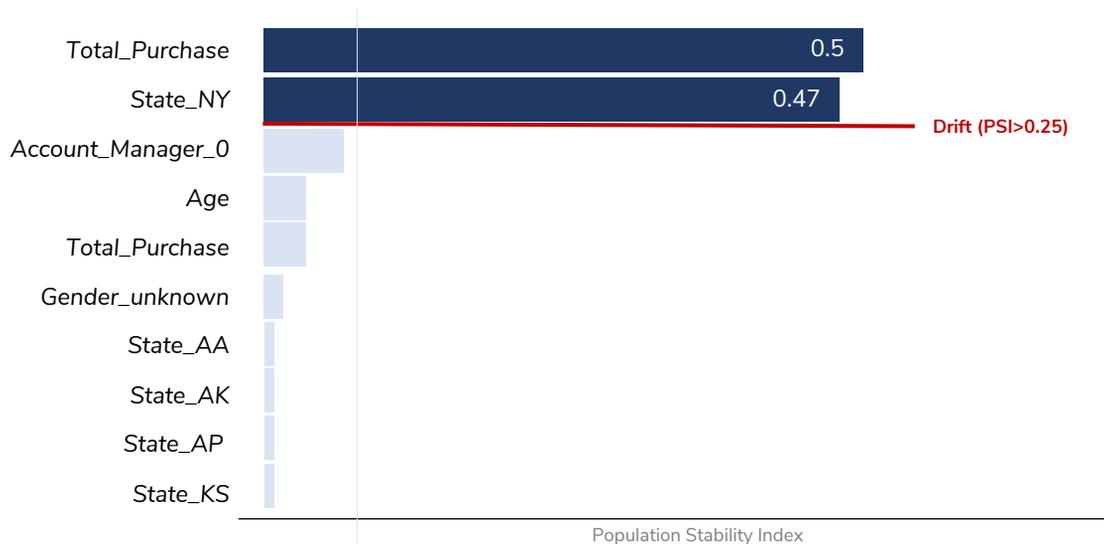
What is data drift?

Data drift is the significant change in data distribution, compared to the data used to train the model. Data drift can be caused by the evolution of business processes, major industry events altering data trend, etc. It does not affect the formatting of the data, but the data at its core and what it represents.

Why is data drift important?

Data drift leads to the degradation of a model's performance on new data. In production data are now significantly different to the data used to train the model, resulting in less accurate predictions. The model «does not know» how to correctly predict on these new data, since its training did not include comparable observations.

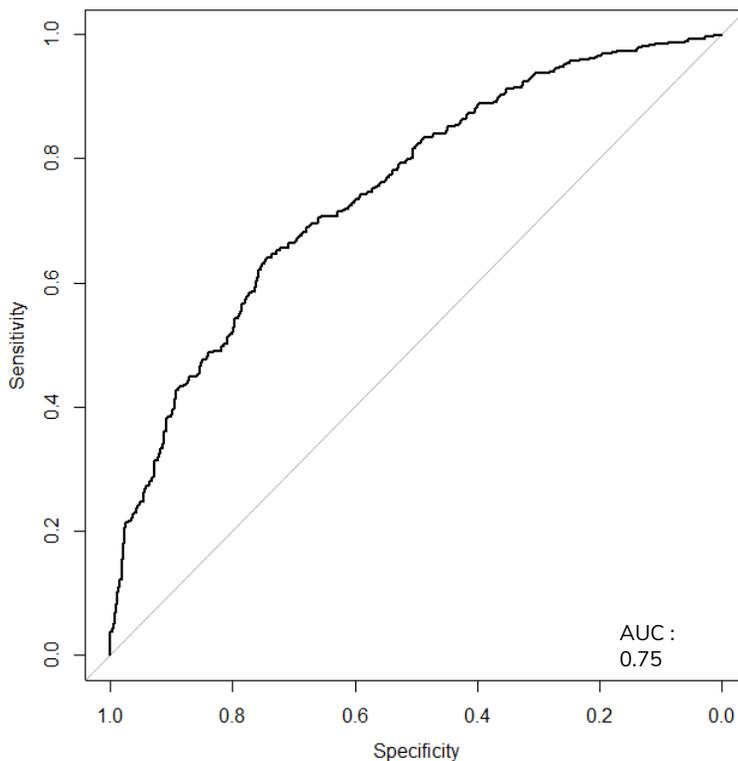
Population Stability Index Test Details



How Snitch evaluates data drift

Snitch uses a series of indicators and metrics to evaluate the magnitude of the differences between the production data and the model's training data. This includes the classic population stability index from econometrics and checking if a simple classifier can distinguish the training data from the new operational data.

Production Observations Separable From Training Observations Test Details



A simple classifier's AUC is greater than 0.6, it can therefore distinguish the training data from the production data

How Snitch evaluates data drift

Snitch checks if a simple classifier can distinguish the training data from the production data. If yes, production data distributions do not match training observation distributions. This means that production data is not similar to the training observations and it could affect the model's ability to produce correct predictions.

Compare model performance on Training and Testing**FAIL****Observation**

- Model's performance on Training is 10% better than on Testing.

Recommendation

- Snitch recommends to re-train model using techniques to avoid overfitting by :
 - Adding dropout layers
 - Adding a regularization method
 - Training model with cross-validation
 - Set an Early Stop when training the model

Residual analysis for Training and Testing**FAIL****Observation**

- For class 0, residual for Training is lower than residual for Testing

Recommendation

- Snitch recommends to re-train the model using techniques to avoid overfitting by :
 - Generating synthetic observations for Training
 - Set an Early Stop when training the model

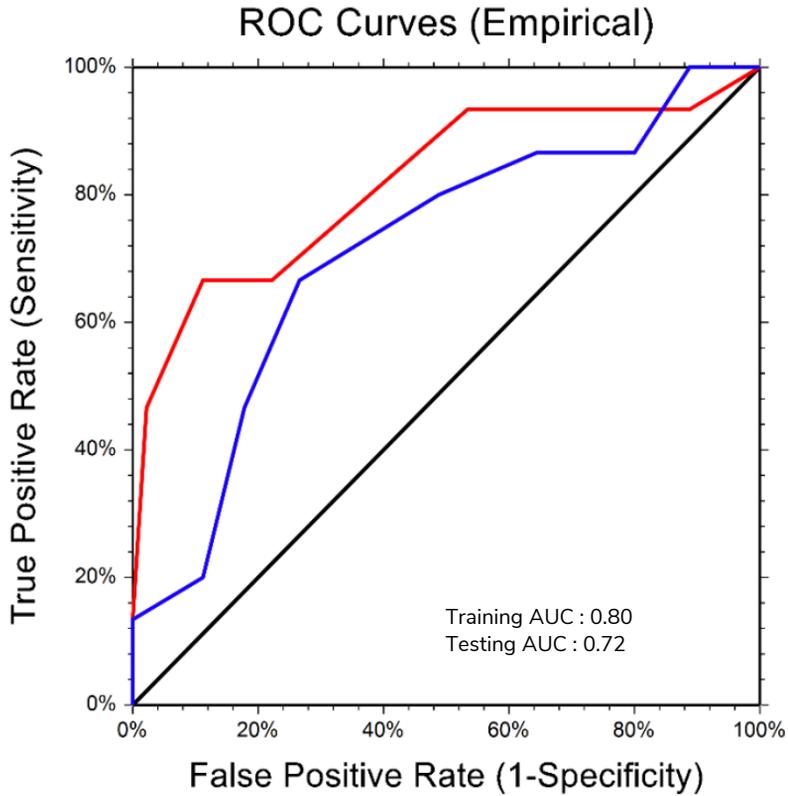
What is overfitting?

A model is overfit on the training data when it is not able to generalize and predict on new, never seen before data. The training data is modeled too closely, often due to overly complex models.

Why is overfitting important?

Overfit models cannot adequately perform on real-world production data. Performance on the training data can be exceptional, but quality of predictions on never seen before data will be significantly poorer.

Compare model performance on Training and Testing

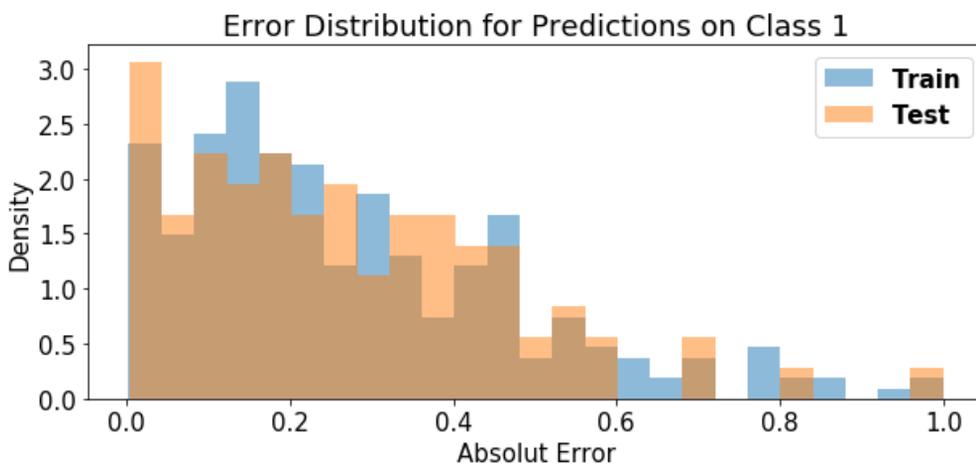
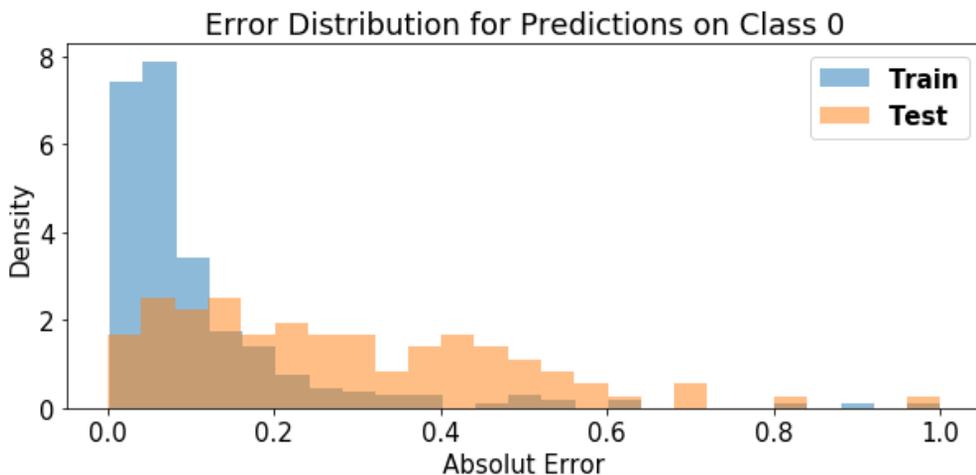


Model performance on Training dataset is significantly greater (+8%) than model performance on Testing dataset, indicating that the model has overfit.

How Snitch evaluates overfitting

Snitch compares the model's performance on training data and testing data. If the performance is significantly better on training data, then the model has overfit.

Compare model performance on Training and Testing



How Snitch evaluates overfitting

Snitch compares the model's performance on training data and testing data. It then progressively adds dropout to the model and generates predictions. By sampling enough times, it is possible to estimate the distribution of predictions and therefore calculate a confidence interval. Overfit models will have a larger confidence interval on new data than on the well-known training data.

Model predict_customer_churn	Characteristics - 540 Training observations - 180 Testing observations - 180 Production observations	- 74 Features - Classification model - Sequential architecture
Version 1.1.2		

Model file integrity

File: predict_customer_churn.h5

Size: 52.1 MB

Hash (SHA3-512):

eb3fd4cd3b7a6dd6e6077f058ea5147ff8c9010aeaba21234bce30a602d1e964c21509c5a3fbbd70ce0bb81ae2c2d96529a9392f5584b56f13def7c6f8dabe94

Snitch seal of authenticity

This document was generated automatically by Snitch on June 6th, 2020 and timestamped through Digicert RFC3161 server <http://timestamp.digicert.com>.



Digitally signed by Snitch

DN: cn=Snitch, o=Moov AI Labs, email=info@snitch.ai, c=CA

Reason: This document is certified authentic.

Date: 2020.06.06 08:55:32 -04'00'