Credit Score Classifier

Group 4: Jimmy Nguyen, Kyle Jow, Kyle Pickle, Jacob Tuttle, Kris Wong

Background

- Credit scoring necessary to mitigate lending risk
- Banks use it to predict whether a loan will be paid back in full or not

- One of the oldest implementations of ML (1950s)
- 5 C's approach ➪ Decision Trees ➪ Log. Reg. ➪ ??
- Performance vs Transparency

- We use credit Score data from Kaggle
- Features based on FICO credit standards
- Used to classify credit score into:
  - Good
  - Standard
  - Poor

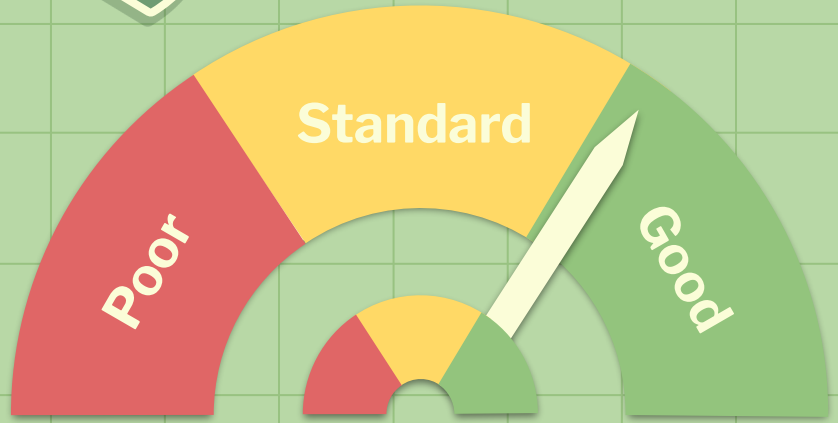

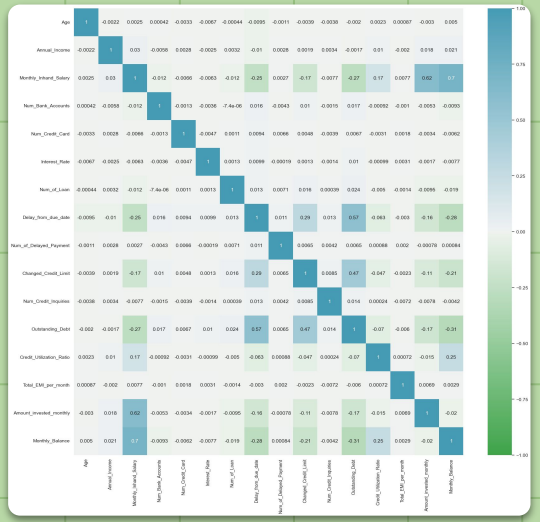

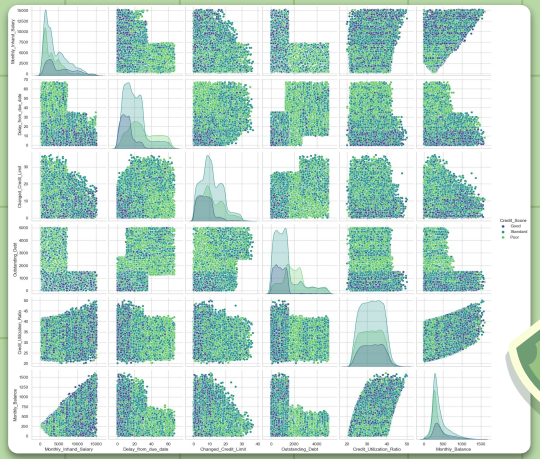

Exploratory Analysis

- Fairly imbalanced dataset
  - Not enough ‘Good’ classifications
- Outliers and unnecessary features

- Decided to use different models to test the dataset:
  - SVM Linear Kernel
  - SVM RBF Kernel
  - Multinomial Logistic Regression

- Kaggle dataset had 27 features
  - Cut down these features to 5
- 80:20 split

Results

Logistic Regression

	precision	recall	f1-score	support
Good	0.50	0.14	0.22	2801
Poor	0.58	0.39	0.47	4560
Standard	0.57	0.81	0.67	8213
accuracy			0.57	15574
macro avg	0.55	0.45	0.45	15574
weighted avg	0.56	0.57	0.53	15574

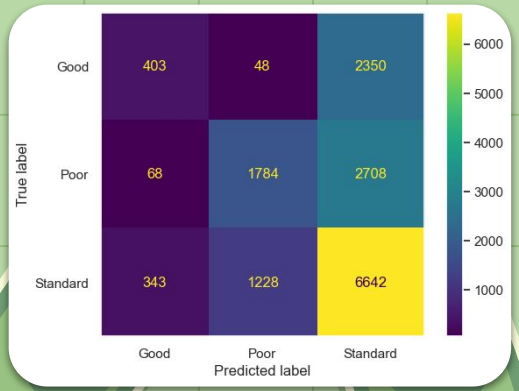

SVM Linear Kernel

	precision	recall	f1-score	support
Good	0.00	0.00	0.00	2801
Poor	0.59	0.44	0.50	4560
Standard	0.56	0.84	0.68	8213
accuracy			0.57	15574
macro avg	0.39	0.43	0.39	15574
weighted avg	0.47	0.57	0.50	15574

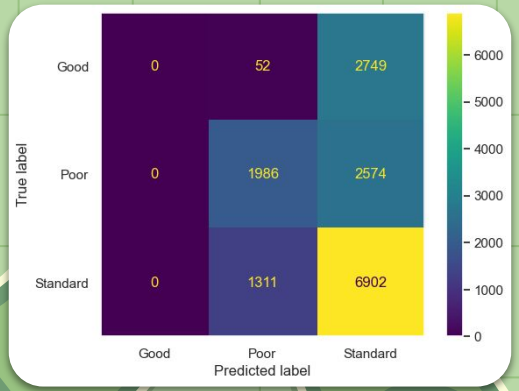

SVM RBF Kernel

	precision	recall	f1-score	support
Good	0.59	0.06	0.10	2801
Poor	0.66	0.57	0.61	4560
Standard	0.60	0.83	0.70	8213
accuracy			0.62	15574
macro avg	0.62	0.49	0.47	15574
weighted avg	0.62	0.62	0.57	15574

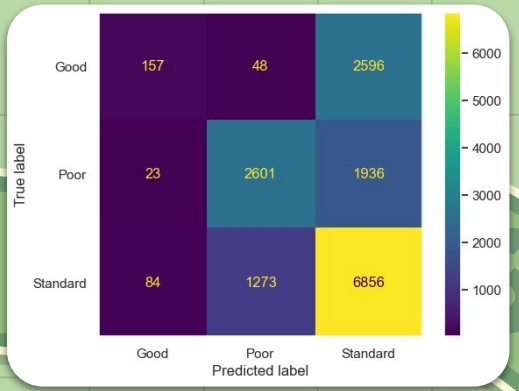

Results - Cont'd

Logistic Regression

- Outliers

	precision	recall	f1-score	support
Good	0.52	0.28	0.37	2660
Poor	0.62	0.40	0.49	3715
Standard	0.61	0.81	0.69	7506
accuracy			0.60	13881
macro avg	0.58	0.50	0.52	13881
weighted avg	0.59	0.60	0.58	13881

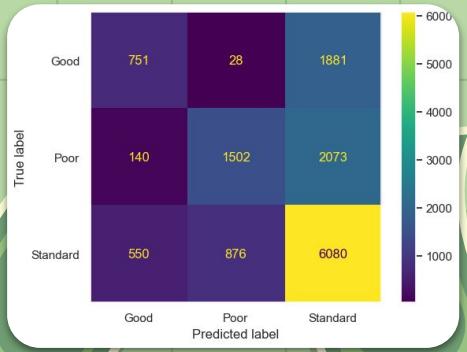

SVM Linear Kernel

+ Oversampling - Outliers

	precision	recall	f1-score	support
Good	0.40	0.75	0.52	2660
Poor	0.62	0.68	0.65	3715
Standard	0.75	0.49	0.59	7506
accuracy			0.59	13881
macro avg	0.59	0.64	0.59	13881
weighted avg	0.65	0.59	0.59	13881

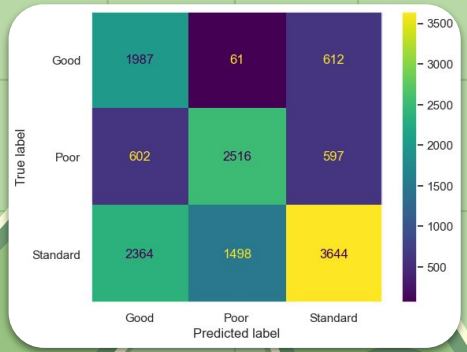

SVM RBF Kernel

+ Oversampling - Outliers

	precision	recall	f1-score	support
Good	0.40	0.75	0.52	2660
Poor	0.62	0.68	0.65	3715
Standard	0.75	0.48	0.59	7506
accuracy			0.59	13881
macro avg	0.59	0.64	0.59	13881
weighted avg	0.65	0.59	0.59	13881

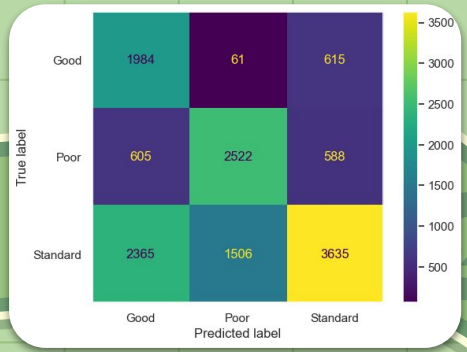

Conclusion & Analysis

- SVM with RBF kernel yielded the highest accuracy
  - Highest recall for standard credit scores
  - Highest Precision for poor credit scores
- Very low performance in all models when predicting high credit scores
  - Possibly increased by sampling an equal amount of each credit score

- Removed outliers in logistic model:
  - Slightly better at predicting scores, especially ‘Good’
- Oversampling and removing outliers in SVM models:
  - Greatly improved predictions of edge ‘Poor’ and ‘Good’ credit scores
  - ...at the cost of ‘Standard’

Demo

https://ecs171-project.streamlit.app/