Drawing the Line: Classifying Risky Mortgage Applicants with ML

Alex Arrieta | Josh Blank
Brendan Callender | Sophia Chung
Martin Hsu

Project Overview

Mortgage Risk
Management

Credit Default
Classification

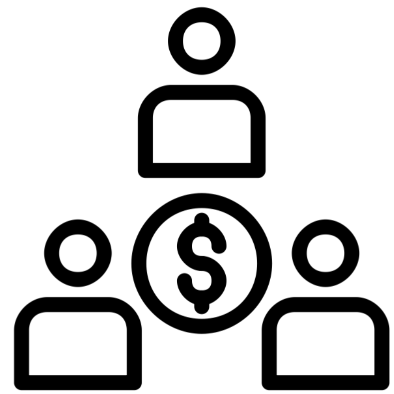

The “Ideal”
Applicant

Agenda

1. Data Context
2. Variable Selection
3. Model Selection
4. Final Model
5. Model Performance
6. Model Fairness
7. Key Findings
8. Recommendations

Data Context

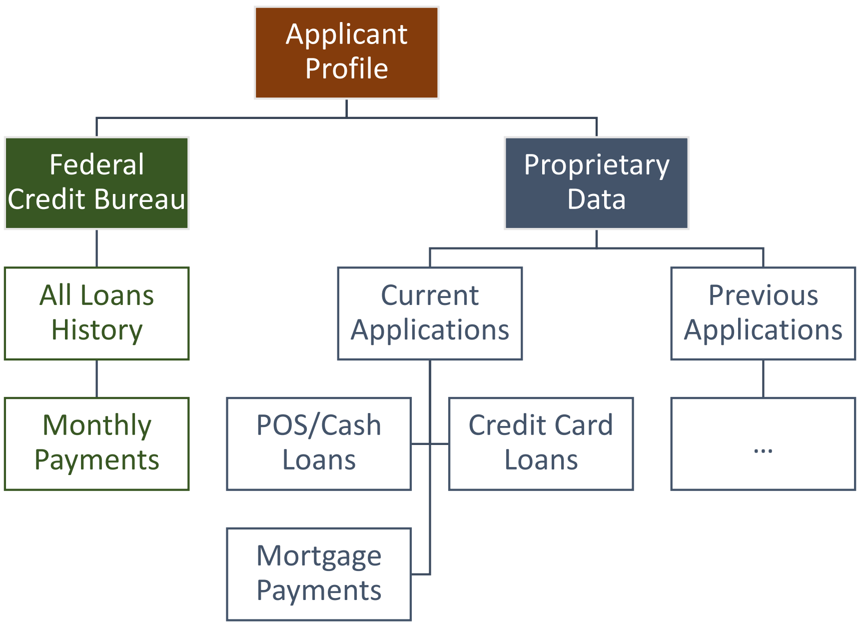

- Classification Target: Did **applicant default**? (Y/N)
- Case Study: **Home Credit Group©**
  - Includes ground truth
- Each row in primary dataset is a loan **application**

Variable Selection

Prediction Target

Applicant Profile

Application ID	Defaulted?	Education	Gender	Owns Car	Income	...	Credit Amount
100002	**1 (Yes)**	Secondary	M	1 (Yes)	$202,500	...	$406,598

Proprietary Mortgage History

Prev Approved Loans	Prev Canceled Loans	No. Late Payments	...	No. Completed Contracts
1	0	0	...	0

Federal Credit Bureau Loan History

Mortgage Total Days Overdue	Other Loan Total Days Overdue	Mortgage Total Amount Overdue	Other Loan Total Amount Overdue
0	0	0	0

Model Selection

Logistic Regression Classifier

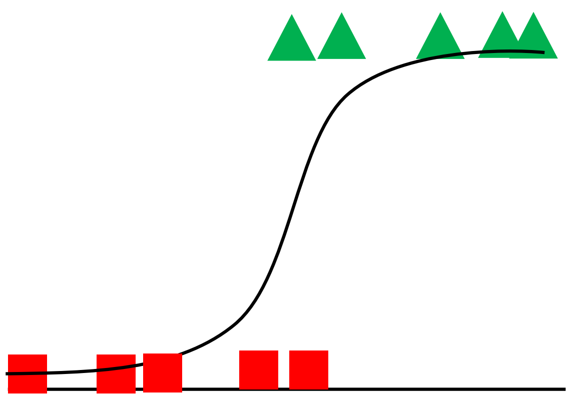

Linear Relationships

Support Vector Classifier

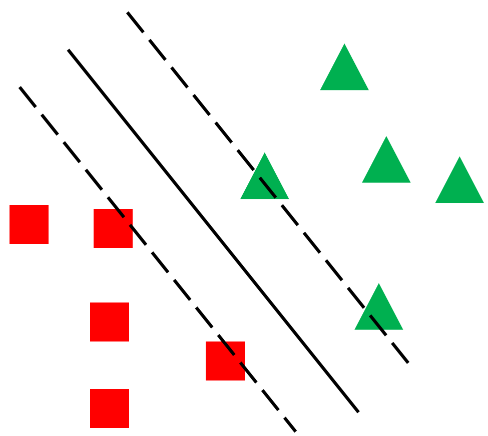

Maximal Separation

Linear Discriminant Analysis Classifier

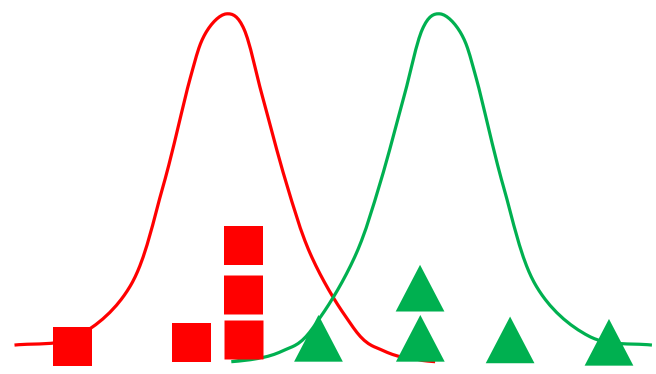

Highest Likelihood

Model Selection

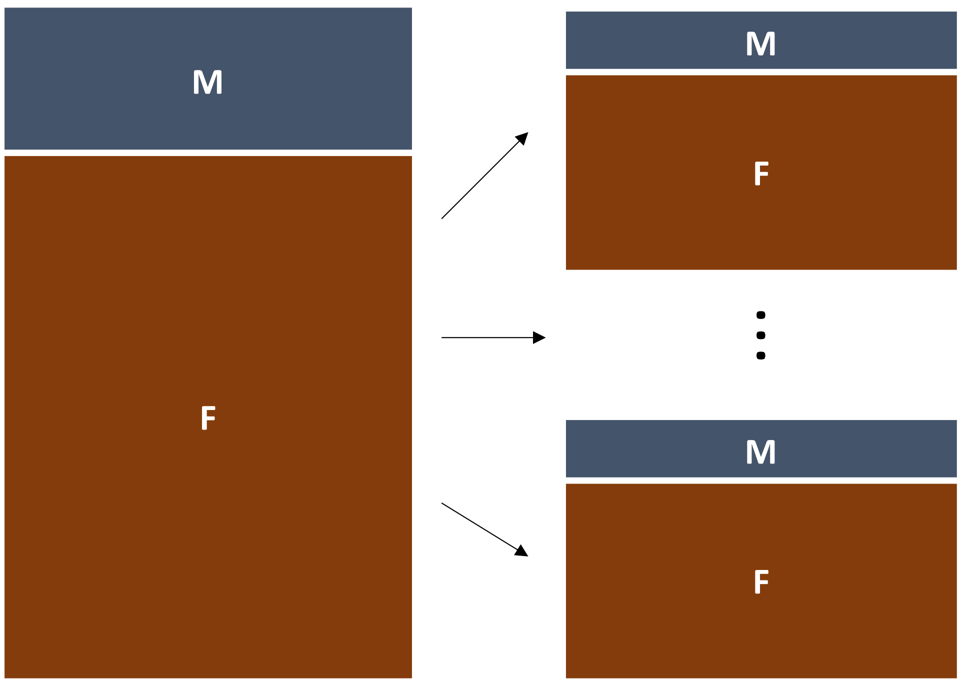

Testing in a way that **enhances predictive power** and **fairness**

- **Splitting:** "Training" and "testing" models on different samples to simulate new data
- **Stratified sampling:** Preserving representation

Final Model

Applicant Profile Variables
Age (Days)
Car Ownership
Education
Income Type
Loan Defaults in Social Circle
Region Rating

Proprietary Application Variables
Number of On-Time Payments
Number of Payments
Prior Application Approval/Denial

Linear Discriminant Analysis (LDA)

Chosen by accuracy and fairness metrics

Can output both **risk score** or "**probability**" of default

High default risk: > 0.2 prob

Low default risk: < 0.2 prob

Model Performance

Accuracy	Precision	Recall
(out of 100%)	(out of 100%)	(out of 100%)
61%	**47%**	**62%**
…of applications **correctly predicted** overall	...of applications **predicted** to default, **actually** defaulted	…of applications **actually** defaulted were **predicted** to default

Model Performance

F1 Score	ROC-AUC Score
(out of 1.00)	(out of 1.00)
0.54	**0.59**

Moderate overall performance in both **precision** and **recall**

Fairly accurately **distinguishes** between defaults and non-defaults

Model Fairness

Women are correctly predicted as **defaulting** at a rate of:

True Pos. Equalized Odds
(Ratio)
1.01

…**times** that of men.

Women are correctly predicted as **not defaulting** at a rate of:

True Neg. Equalized Odds
(Ratio)
0.95

…**times** that of men.

Key Findings

The “Ideal” Applicant

Variables that lower risk score

Key Findings

Variables that increase risk score

The High-Risk Applicant

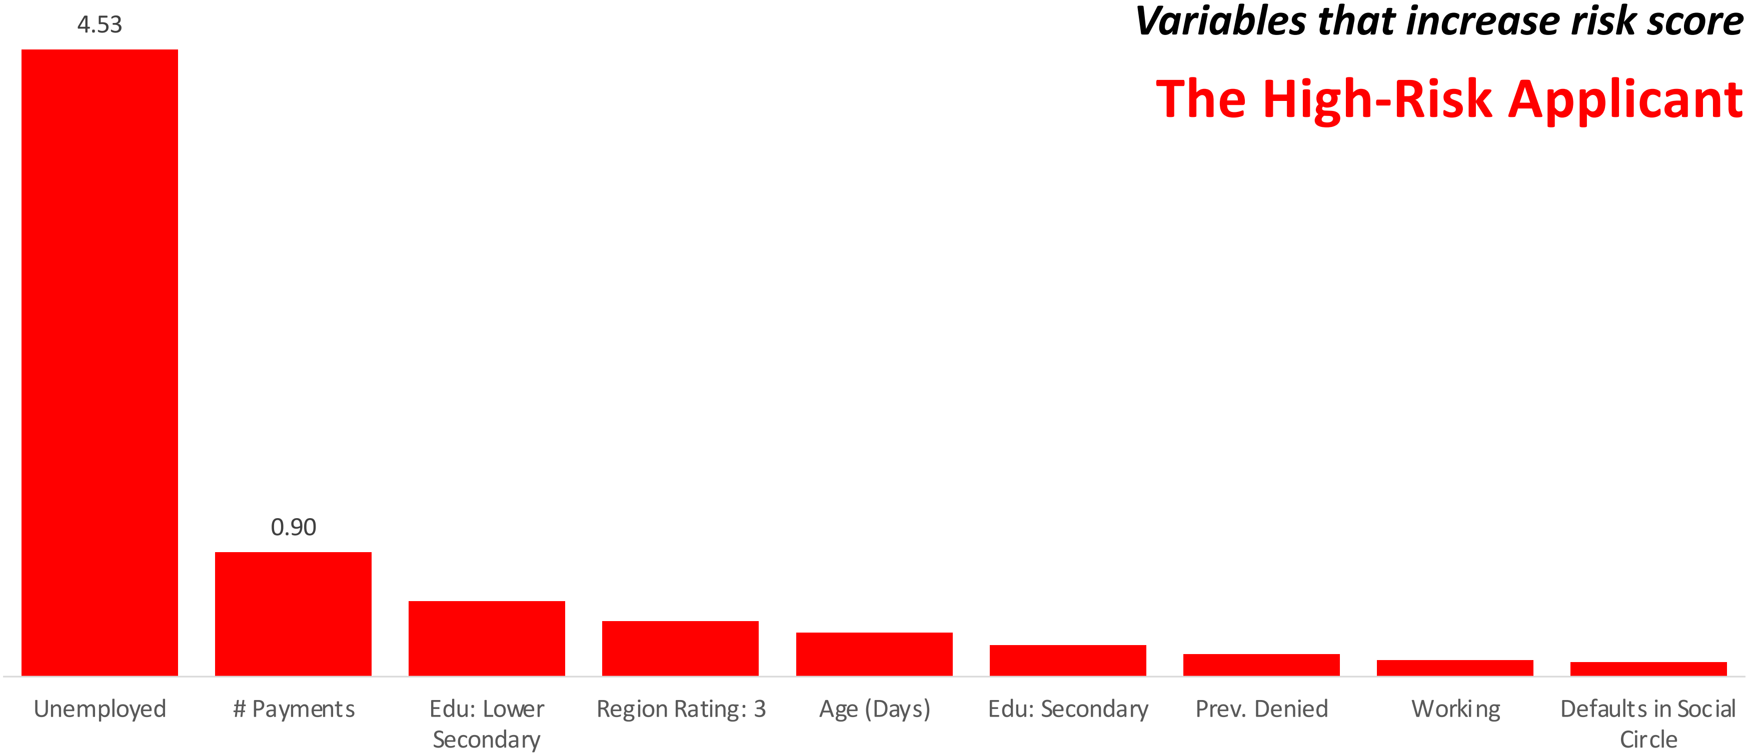

Recommendations

The “Ideal” Applicant is

Educated

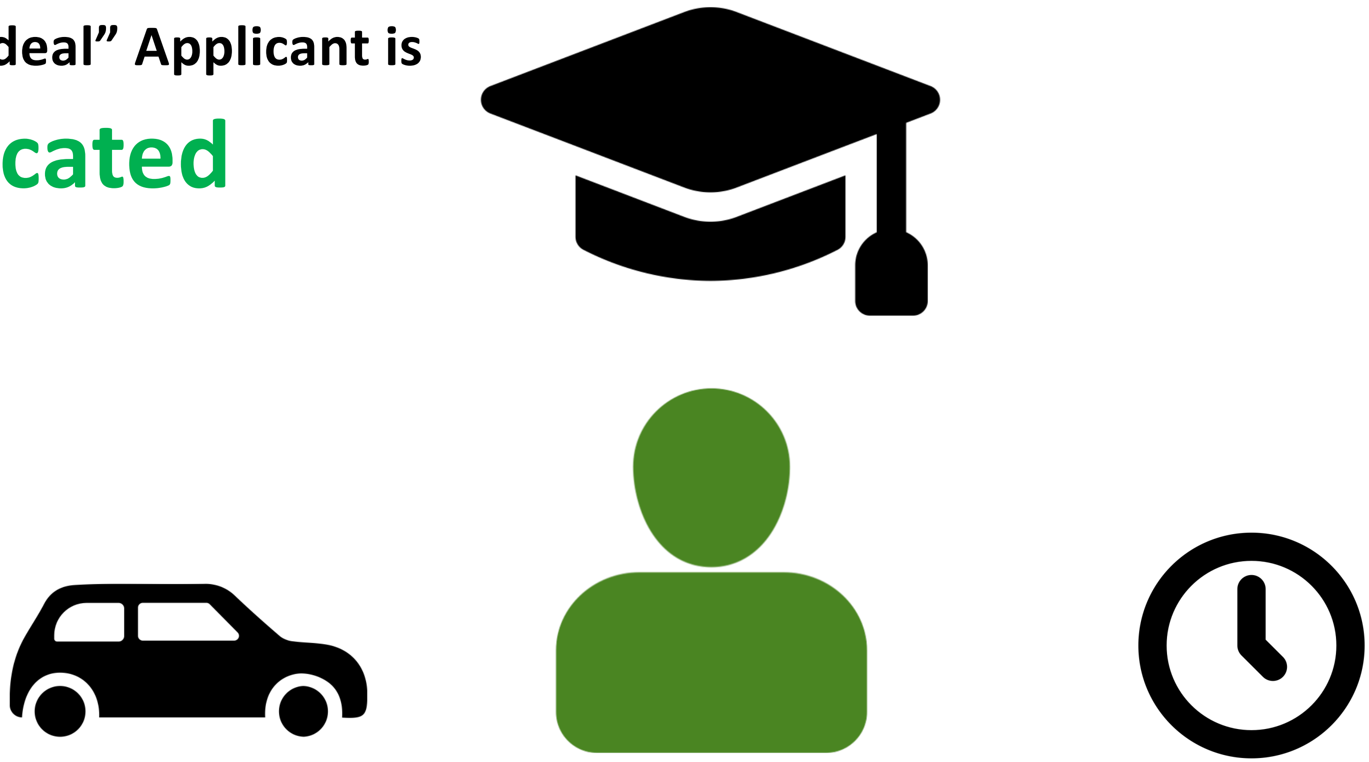

Recommendations

The High-Risk Applicant is

Unemployed

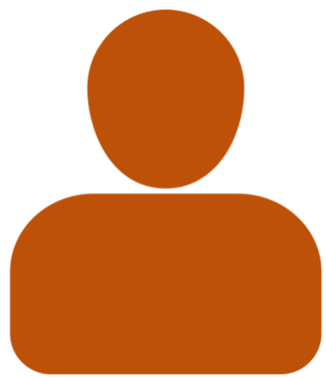

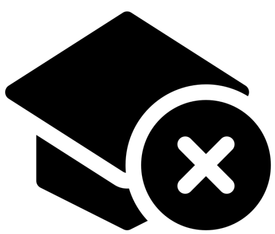

Alex Arrieta | ajarriet@calpoly.edu
Josh Blank | jwblank@calpoly.edu
Brendan Callender | bscallen@calpoly.edu
Sophia Chung | spchung@calpoly.edu
Martin Hsu | mshsu@calpoly.edu