Quantitative Risk – Decision-Making Models & The Use of Advanced Estimation Techniques [Machine Learning]

August 2023

Background

Modelling techniques for bank risk management have always been an important element, with greater emphasis during the last two decades. Well established risk quantification methods are used by banking institutions within their capital calculation, provisioning, forecasting and stress testing, pricing and decision making. The significant improvement in data processing and computational capabilities has resulted in an increasing industry trend to use more advanced techniques in risk identification and quantification. The trend has been strongest in the area of decision making (non-regulatory) models. However, recent publications on regulatory models from the EBA (European Banking Authority) and PRA (Prudential Regulation Authority) are developing advanced method use cases. An increased focus on Artificial Intelligence (AI) and machine learning (ML) methods highlights the need for bank risk management to understand the capabilities of advanced modelling techniques.

Within Credit Risk, a role for ML models is becoming more and more relevant across several areas. Examples include:

- Deciding how to categorise or rank order loans or borrowers;
- Identifying the cohort of loans that are most vulnerable;
- Designing efficient credit sanctioning and review processes and
- Identifying emerging risks in a dynamic way.

Having the ability to process data efficiently and help institutions to make better informed decisions in a timely manner represents a few benefits of advanced modelling techniques.

We explore **three different methodologies to measure the discriminatory power** between good and bad borrowers using a credit card portfolio dataset.

Executive Summary

Our focus in this paper is to develop decision making models using a range of advanced machine learning techniques. We explore three different methodologies to measure the discriminatory power between good and bad borrowers using a credit card portfolio dataset. The main hypothesis is that advanced modelling techniques lead to more efficient estimates and higher discriminatory power.

The first section reviews the applicable regulatory information. Recent regulatory publications and initiatives acknowledge a need to address advanced modelling techniques use cases in bank risk management. A focus on the regulatory capital models shows a relatively conservative regulatory position. The second section provides the definitions on decision making and machine learning principles. In the third section, three models with varying complexity levels are applied to the credit card dataset.
An assessment is performed on the ability to discriminate borrowers repayment status between good (performing) and bad (nonperforming). Model output confirms the hypothesis that applying more advanced ML techniques significantly improves model discriminatory power, leading to more efficient rank ordering.
A final section provides a benefit analysis (pros and cons) of using ML techniques in this area.

Relevant Regulation

During the last few years, competent supervisory and regulatory institutions (European Central Bank (ECB), European Banking Authority (EBA) and Prudential Regulation Authority (PRA)) have started to address the use cases of applicable machine learning techniques to bank risk management frameworks.

For regulatory models that impact capital (regulatory or economic) or accounting guidance on provisioning, a conservative mind-set remains in place. Consensus shows that application of advanced ML techniques within the banking space is at an early stage. As data availability and complexity increases exponentially, a clear trend emerges for institutions to implement advanced analytics in their different risk frameworks.

EBA Thoughts To Advanced ML Modelling Techniques

In November 2021, the EBA published a discussion paper on machine learning used in the context of IRB models[1] to calculate regulatory capital for credit risk.

An aim of the discussion paper was to set supervisory expectations on how new sophisticated machine learning models can co-exist with and adhere to the Capital Requirements Regulation [CRR] when used in the context of IRB models.

Feedback was requested on many practical aspects related to the use of ML in the context of IRB, with the aim of providing clarity on supervisory expectations regarding their use.

During August 2023, the EBA published a follow-up report[2] to the previous consultation paper (published in 2021) presenting the main findings and conclusions on the Machine learning in IRB Models topic. A connection was made with the IRB's prudential requirements to other legal frameworks, such as the Artificial Intelligence (AI) Act and the General Data Protection Regulation (GDPR).

The main objective of the Discussion Paper (DP) was to analyse why machine learning and Big Data techniques are being used less in credit risk determination of capital requirements, given the continued exponential increase of data availability and data analysis capability in the financial sector. The DP also explores potential issues related to the compliance of these techniques with relevant regulatory guidance on the use of IRB models.

Following the 17 proposed questions on the DP, the EBA received 14 responses. Institutions agreed that issues regarding implementation of ML techniques for IRB Models require time to mature.

Four main pillars were identified by the EBA for Institutions to focus on:

1. Data Management
2. Technological Infrastructure Enhancement
3. Organisation And Governance Towards Regulatory Compliance
4. Analytics Methodology

1 EBA/DP/2021/04 - EBA Discussion Paper on Machine Learning for IRB Models
2 EBA/REP/2023/28 - Follow-Up Report from the consultation on the Discussion Paper on Machine Learning for IRB Models

PRA Thoughts To Advanced ML Modelling Techniques

PRA, in its paper on Artificial intelligence and machine learning, [DP5/22] indicated that the use of the advanced modelling techniques could enable firms to:

- Offer better products and services to consumers;
- Improve operational efficiency;
- Increase revenue; and
- Drive innovation,

all of which may lead to better outcomes for consumers, firms, financial markets, and the wider economy. According to the PRA, there are three stages in the AI lifecycle:

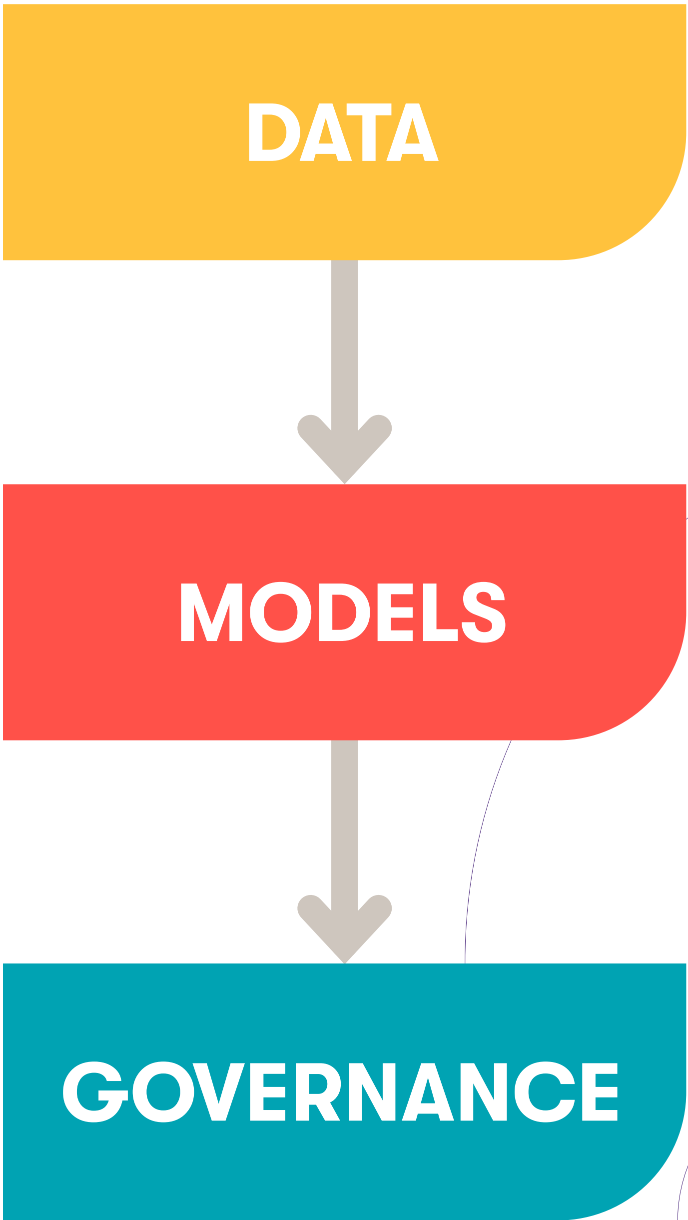

It is important to realise that interconnected risks still remain in place. Elements at the data level can feed into the model level. There is potential to raise broader challenges at the level of the firm and its overall governance of AI systems.

The publication on Model Risk Management Principles For Banks [PS6/23] outlines eleven responses to the question on MRM for AI/ML models. Overall, respondents supported the PRA's proposals to raise the standard of MRM practices and recognised the need to manage risks posed by models that have a material impact on business decisions.

Respondents were in broad agreement that the principles are sufficient to identify, manage, monitor, and control the risks associated with AI/ML models. The main areas highlighted by respondents were as follows:

- AI/ML systems often span multiple functional areas, including data, models, and technology;
- Explaining how an AI/ML model works and how it produces its outputs can be challenging;
- Some AI/ML models are dynamic by design i.e. they can change and/or recalibrate frequently;
- Monitoring of model performance becomes increasingly important as AI/ML model complexity increases; and
- Several respondents pointed out that the use of AI/ML models can raise ethical challenges, including fairness and bias.

Decision Making Models and Machine Learning

According to the IIF 2019 report, machine learning is widely used in credit risk, particularly in credit decisions/pricing, credit monitoring, collections, restructuring, and recovery. However, ML techniques are generally avoided in regulatory areas like capital requirements for credit risk, stress testing, and provisioning, due to their complexity and challenges in interpretation and explanation.
In the context of IRB models in compliance with CRR requirements, examples of where ML techniques are currently used include:

Model Validation

ML is used to develop model challengers to serve as a benchmark to standard models used for calculating capital requirements.

Data Improvements

ML algorithms can efficiently handle large datasets, identify data patterns, and handle missing or noisy data, resulting in improved model data quality.

Variable Selection

ML algorithms can aid in detecting relevant explanatory variables via feature engineering optimisation within large datasets.

Risk Differentiation Of The Probability Of Default (Pd) Model

ML allows for PD grade movements using text mining algorithms. Qualitative data can be analysed using text mining techniques to extract valuable insights that can impact credit risk assessments.

Decision-Making Models

Refer to the systematic frameworks or methods used to assess and evaluate credit risks associated with lending or extending credit to individuals or businesses. Allow financial institutions and lenders to make informed decisions about whether to:

- approve or deny credit applications;
- set appropriate credit limits;
- determine interest rates; and
- manage overall credit risk exposure.

Credit Risk Decision-Making Models

Typically involve analysing various factors and information to assess the borrowers' likelihood of default or delinquency. Use statistical techniques and predictive analytics to:

- estimate the probability of default;
- evaluate the creditworthiness of borrowers; and
- determine appropriate risk mitigation measures.

The Use Of Ml In Decision-Making Modelling

Has emerged as a transformative force in the realm of decision-making. Offering unprecedented capabilities to build sophisticated models that can analyse complex data, predict outcomes, and optimise choices across a wide range of domains. By harnessing the power of ML algorithms, AI can efficiently process vast amounts of information, identify patterns, and generate valuable insights that human decision-makers might overlook. Through this synergy of human expertise and AI-driven analysis, decision-making models can be crafted with greater accuracy, efficiency, and adaptability, empowering businesses, governments, and individuals to make well-informed choices in an increasingly dynamic and data-driven world.

Moreover, ML-driven decision-making models have an advantage of adaptability with continuous improvement/development. Learning from new data and experiences, the technology can dynamically make predictions and recommendations. This ensures that decisions remain relevant and up-to-date in an ever-changing environment.
Such adaptability is especially crucial in industries where circumstances fluctuate rapidly, such as finance, healthcare and marketing.

Machine Learning Theory

Types Of Machine Learning Algorithms

Supervised Learning

- The algorithm learns rules to building a model from a training dataset that contains a target variable.
- Model rules enable prediction of the outcome variable when presented with new input data.

Unsupervised Learning

- The algorithm learns from a training dataset which has no target variable.
- The goal of the algorithm is to understand the distribution of the data by identifying interpretable patterns, associations, and descriptive properties.

Reinforcement Learning

- Instead of relying on a training dataset, the algorithm learns through interactions with an environment, adapting its behaviour through trial and error.
- The algorithm goal is to perform a specific task effectively by optimising its actions based on the feedback received from the environment.

How A Machine Learns

Machine learning combines tools from both statistics and computer science to generate efficient models. The aim is to make predictions through learning e.g., using large-scale datasets. A fascinating feature of ML is the way the models “learn” or “train”. This provides a clear distinction between traditional programming and ML. Traditional programming incorporates a mechanical process via programming to use data and “rules” to produce results. However, in machine learning, the system develops a model using the discovered pattern in the dataset. Such models can be utilised to hyopthesise real-world phenomena.

Theoretical Components of ML

Background knowledge relevant for machine learning algorithms relates to the mathematical components of **optimisation, function approximation** and **probability and statistics.**

- **Optimisation** is essential for determining the best performing model in a class of models. It is usually completed by fitting a function to the observed dataset. It entails finding input parameters that either minimise or maximise the objective function of the intended model. This can involve an iterative process and is one of the most important theoretical concepts in machine learning. Commonly utilised optimisation algorithms include:
  - gradient descent;
  - stochastic gradient descent; and
  - minibatch stochastic gradient descent.
- **Function approximation** entails understanding how to examine variable dependency in a dataset. Using a learning algorithm, a relationship is represented with a mathematical function or mapping. Neural networks represent a model that is utilised as a function approximator in classification problems, by mapping the inputs to the class labels in a dataset and combined with other algorithms to adjust weights.
- **Probability and statistics** are essential components in the theory of machine learning, as they help to measure the uncertainty in making futuristic predictions. A knowledge of stochastic processes, with their binding probability rules, impacts the ability to make such predictions on future observations.

Model Build Methodology

The continued growth of big data across all sectors encourages responsible use of such resources. Exploring research approaches with a sustainability lens becomes more important. In recent years, machine learning methodologies seek to provide solutions. Different algorithms can help to carry out better analysis by generating more accurate and precise results with reduced error levels.

Machine learning algorithms have been increasingly utilised across finance, ranging from fraud detection to credit risk modelling, when plausible amounts of data are available. Our research has utilised the three learning algorithms of Logistic Regression, Random Forest and Extreme Gradient Boosting. Results are compared to determine the best-performing model.

Logistic Regression

Logistic Regression is a **linear** classification algorithm used as supervised learning where the target (dependent) variable or label is dichotomous; that is, having two outcomes, such as good/bad, yes/no. It can provide an initial benchmark for modelling datasets with such dependent variables when it contains minimal ambiguity. For modelling the credit risk of potential clients, the algorithm minimises the empirical risk (mean squared error) by maximising the likelihood of the training dataset.

The Logistic Regression for the dichotomous dependent variable seeks to evaluate the probability given as:

$$p(x) = \mathbb{P}(Y = 1 | X = x)$$

Thus, formally, the Logistic Regression model is represented as

$$\log\left(\frac{p(x)}{1 - p(x)}\right) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_k x_k$$

Where $\beta_0, \ldots, \beta_k$ are parameters and $x_1, \ldots, x_k$ predictor variables.

Moreover, Logistic Regression utilises the logistic function. Ensuring that the linear combination of inputs from the dataset returns a value in the range 0 to 1. Displayed in the formula as:

$$\sigma(t) = \frac{1}{1 + e^{-t}}$$

Where $t = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_k x_k$

Logistic Regression is a widely used method in classification tasks. It is both **easy to implement** and serves as a **quick initial benchmark**

Random Forest

Single machine learning algorithms tend to overfit the dataset, which can impact the model performance metrics. Hence the inclusion of ensemble learning methods. This approach entails combining different learning algorithms to make predictions. A Random Forest is an ensemble method involving a combination of multiple decision trees to generate improved outputs.

A Random Forest model is a classifier composed of multiple trees. If the number of decision trees is k, then the collection of tree classifiers is defined as

$$h(x, \Theta_k), k = 1, \ldots$$

where each Θk are independently and identically distributed random vectors, with each tree casting a single vote for the modal class at input x.

Random Forests tend to achieve **higher accuracy compared to individual decision trees, especially on complex and high-dimensional datasets.** They are less prone to overfitting due to the averaging effect of multiple trees, which helps improve generalisation to unseen data

Extreme Gradient Boosting (XGBoost)

Extreme Gradient Boosting, or simply XGBoost, is a widely recognised ensemble technique among gradient boosting algorithms. It revolves around the iterative amalgamation of weak learners within the model. XGBoost has gained prominence due to its rapid execution, straightforward integration, and impressive capabilities when handling vast datasets. Notably, its objective function incorporates L1 and L2 regularisation components that contribute to enhanced performance and the mitigation of overfitting concerns.

This is defined as:

$$\mathcal{L}(\phi) = \sum_i l(\hat{y}_i, y_i) + \sum_k \Omega(f_k)$$

where the first term and the second terms are the loss function and regularisation, respectively.

Regularisation

The combination of regularisation, tree pruning, and column subsampling helps prevent overfitting and improves model generalisation. Similarly, XGBoost can handle various data types and works well with both numerical and categorical features.

Hypothesis

Ensemble methods will outperform linear methods

Model Performance Evaluation

The model performance is evaluated using two distinct approaches. The first approach is using confusion matrix that compares the model's predicted class labels with the true class labels from the actual data. The second approach relays on the ROC curve, the graphical representation that shows the trade-off between the true positive rate (Sensitivity) and the false positive rate (1 - Specificity) for different classification thresholds. Both approaches indicate significant improvement in the model performance using advanced machine learning techniques.

Approach 1 – Confusion Matrix

A **confusion matrix** compares the model's predicted class labels with the true class labels from the actual data. The predicted class labels during this process are distinguished using a 0.5 threshold to achieve the binary classification noted below. In binary classification problems, the matrix has four components:

1. True Positive (TP): Positive observation and predicted as positive.
2. True Negative (TN): Negative observation and predicted as negative.
3. False Positive (FP): Negative observation but predicted as positive.
4. False Negative (FN): Positive observation but predicted as negative.

	True Label	
Predicted label	**Positive**	**Negative**
Positive	True Positive	False Positive
Negative	False Negative	True Negative

Once the confusion matrix is obtained, various **performance metrics** can be derived to assess the model's effectiveness:

Accuracy: Measures the total correct classification

$$\frac{\text{True Positives} + \text{True Negatives}}{\text{Total Instances}}$$

Sensitivity: The ratio of true positives to all the positives.

$$\frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$

Specificity: The ratio of true negatives to all the negatives.

$$\frac{\text{True Negatives}}{\text{True Negatives} + \text{False Positives}}$$

Approach 1 - Performance Metrics

The table presents performance metrics for Approach 1 for the three machine learning models. These metrics provide insights into the predictive capabilities and overall generalisation of these models.

By analysing these metrics, we can discern the strengths and weaknesses of each model. A model with high accuracy, sensitivity and specificity score is considered robust. However, a careful consideration of trade-offs between sensitivity and specificity might be needed based on business requirements. A business may require a model that is better at predicting loan defaults over non defaults and vice versa.

The results indicate that both machine learning techniques outperform standard Logistic Regression model with random forests being the strongest performing model.

	Accuracy	Sensitivity	Specificity
Logistic Regression	0.72	0.40	0.85
Random Forest	0.79	0.50	0.91
XGBoost	0.75	0.43	0.89

Logistic Regression

The Logistic Regression model correctly predicted 180/210 non-defaulted loans (true positive) and 36/90 defaulted loans (true negative).

Accuracy:	0.72%
Precision	0.55%
Recall:	0.4%
F1-Score:	0.46%

Logistic Regression Confusion Matrix Metrics

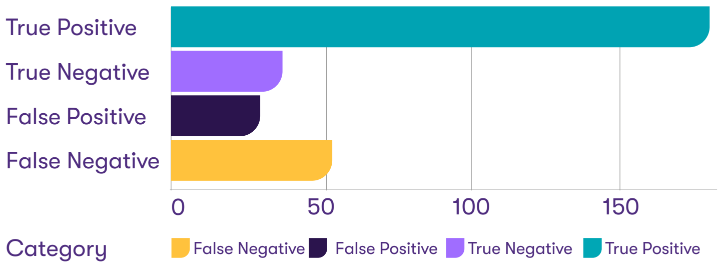

Random Forest

The Random Forest model correctly predicted 192/210 non-defaulted loans (true positive) and 45/90 defaulted loans (true negative). Random Forest models are better at modelling complex non-linear relationships and perform better on larger datasets.

Accuracy:	0.79%
Precision	0.71%
Recall:	0.5%
F1-Score:	0.59%

Random Forest Confusion Matrix Metrics

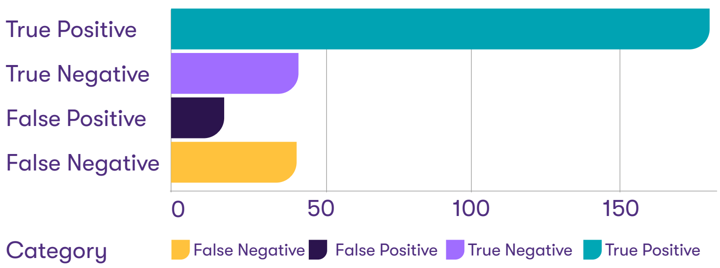

XGBoost

The XGBoost model correctly predicted 187/210 non-defaulted loans (true positive) and 39/90 defaulted loans (true negative). XGBoost models also perform better on larger data sets and hyperparameter tuning allows for better model customisation and performance.

Accuracy:	0.75%
Precision	0.63%
Recall:	0.43%
F1-Score:	0.51

XGBoost Confusion Matrix Metrics

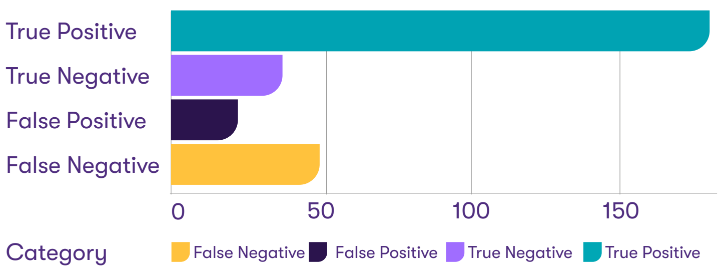

Approach 2 – Receiver Operating Characteristic Curve

Receiver Operating Characteristic Curve

The ROC curve is a graphical representation that shows the trade-off between the true positive rate (Sensitivity) and the false positive rate (1 - Specificity) for different classification thresholds.

Results can be interpreted as follows:

- A good classifier will have an ROC curve that rises sharply towards the top-left corner of the plot, indicating higher TPR and lower FPR; and
- A diagonal line (FPR = TPR) represents random guessing, and any classifier above this line is considered better than random.
- The underlying predicted probability is utilised for the calculation.

Area Under the Curve

The AUC is a single scalar value representing the area under the ROC curve. It provides a measure of a model's ability to distinguish between positive and negative instances.

Result interpretation:

- A perfect classifier has an AUC of 1, indicating that it can perfectly distinguish between positive and negative instances; and
- A random classifier has an AUC of 0.5, equivalent to the diagonal line in the ROC plot.
- The underlying predicted probability is utilised for the calculation.

Gini

The Gini coefficient is used as an alternative evaluation matrix to the AUC, and the two measures are closely related. The Gini coefficient is calculated as twice the area between the ROC curve and the diagonal, or as Gini = 2*AUC-1.

Result interpretation:

- A perfect classifier has a Gini of 1, indicating that it can perfectly distinguish between positive and negative instances; and
- A random classifier has a Gini of 0.5, equivalent to the diagonal line in the ROC plot.
- The underlying predicted probability is utilised for the calculation.

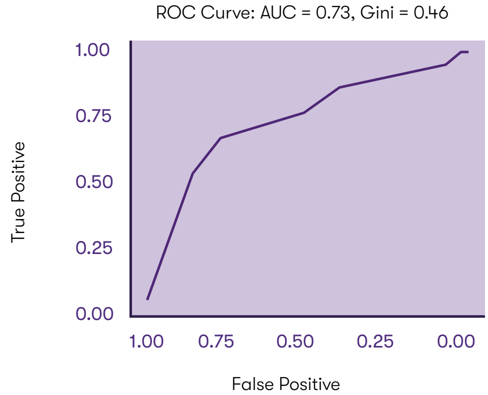

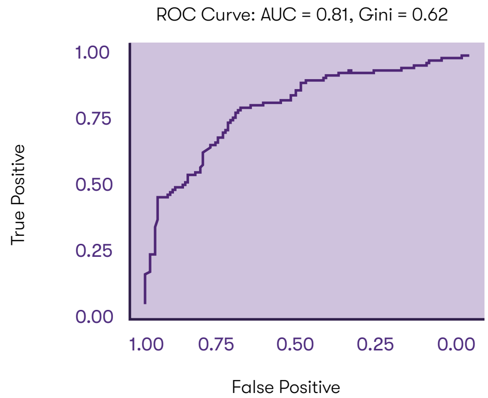

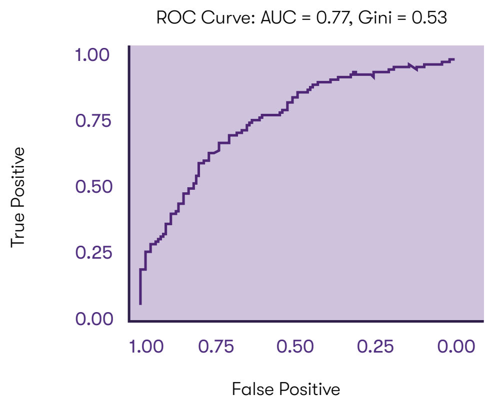

Approach 2 - Performance Metrics

ROC charts and the summary table below presents performance metrics for Approach 2 for the models in scope. These metrics provide insights into the predictive capabilities and overall generalisation of these models. By analysing these metrics, we can discern the strengths and weaknesses of each model. A model with high AUC and GINI score is considered robust. Similarly to the confusion matrix, ROC analysis indicates significantly better performance of machine learning techniques, especially Random Forest which displays much steeper ROC curve with considerably larger area under the curve.

In conclusion, the utilisation of ensemble methods (Random Forest and XGBoost) has consistently demonstrated their superiority over traditional linear model methods. Enhanced model performance can be attributed to the inherent ability of ensemble methods to aggregate individual model strengths and compensate for their weaknesses. Linear models can serve as valuable tools in certain scenarios. The adaptability, robustness, and accuracy offered by ensemble methods make them an indispensable choice for addressing contemporary challenges in decision-making tasks.

Overall, ensemble methods might offer:

- Improved predictive accuracy;
- Ability to handle non-linear relationships;
- Reduced bias
- Competitive advantage supporting enhanced decision making

Challenges & Potential Benefits of Using Machine Learning Techniques

As already outlined, there are both benefits as well as challenges with the use of decision making models that institutions should take into account in the machine learning strategy design.

Benefits

- Automation and Efficiency: Machine learning models can automate and streamline various credit risk assessment processes, reducing the need for manual intervention. Increased efficiency saves time to allow risk analysts to focus on other tasks.
- Better Decision-Making Support: Machine learning can provide valuable insights to decision-makers by identifying relevant risk factors and their impact on credit outcomes. This assists in making more informed and data-driven credit decisions.
- Competitive Advantage: Banks that successfully implement machine learning in their credit risk models gain a competitive advantage by offering more accurate and tailored credit products to customers. These models have the ability to analyse vast amounts of data and identify non-linear patterns that human analysts might overlook. Decisions makers have more information to support data-driven credit decisions, which allows them to assess the creditworthiness of applicants more accurately. This leads to reduced default rates and improves overall portfolio performance.

Challenges

- Data Availability and Quality: Machine learning models require large amounts of high-quality data to learn effectively. Obtaining historical data on credit events and relevant features can be challenging, especially for smaller banks or in emerging markets where credit data may be limited or of lower quality. Limited data can lead to underrepresented patterns and less reliable predictions, potentially hindering the effectiveness of the IRB models.
- Interpretability and Explainability: Machine learning models, particularly complex ones like neural networks, are often considered “black boxes”. A challenge exists to understand and explain their decisions. In regulatory settings such as IRB models, interpretability and explainability are crucial to gain approval and build trust in the models’ results.
- Regulatory Compliance: Machine learning models in IRB frameworks must comply with regulatory guidelines and requirements. These guidelines are designed to ensure that the models are robust, reliable, and meet specific standards for credit risk assessment. Meeting regulatory requirements involves extensive validation, documentation, and reporting processes.

Future Research

- Handling Imbalanced Data: Handling imbalanced data is a crucial aspect of building robust and accurate machine learning models. Should one class proportion be significantly lower than another class, the model can become biased towards the majority class. Future studies could explore methods such as oversampling, undersampling, or generating synthetic data.
    - Oversampling: Involves randomly duplicating instances from the minority class to balance the class distribution. A model would then have more exposure to the minority class during training, potentially leading to better performance.
    - Undersampling: In contrast to oversampling, this method involves randomly removing instances from the majority class to balance the class distribution. By reducing the influence of the majority class, the model may become more sensitive to the minority class.
    - Generating Synthetic Data: Techniques such as the Synthetic Minority Oversampling Technique (SMOTE) create synthetic examples of the minority class by interpolating between existing instances. This seeks to improve the representation of the minority class without simply duplicating existing instances.
- Asynchronous Programming: Asynchronous programming in Shiny apps allows long running and computationally intensive tasks to be executed in the background while the app remains responsive and interactive. Traditional synchronous programming can lead to unresponsive apps, especially when performing complex operations that take time to complete. Asynchronous programming enables tasks to be initiated concurrently, allowing the app to update and display results as soon as each task is completed, even if others are still processing.
- Collaborative Shiny Apps: Enabling collaboration in Shiny apps can significantly enhance decision-making and foster teamwork among users. By integrating real-time data sharing, concurrent editing, and synchronised visualisations, multiple users can work together in the same app, analyse data collectively, and reach informed decisions collaboratively. A collaborative approach fosters a more inclusive decision-making process by reducing communication gaps. This ensures that all relevant perspectives are considered in the credit risk assessments.

Contact

Our team would be delighted to discuss your challenges and opportunities in any aspect of climate risk. Our services are flexible and efficient, designed to facilitate and support your business model. Our highly qualified Quantitative Risk team provides support to financial institutions across the full spectrum of risk measurement and modelling strategies, including the development, deployment and validation of key models and risk measurement methodologies in regulatory capital, stress testing and IRB, IFRS9 and bank risk modelling. Our team has experience implementing machine learning techniques in the context of credit risk modelling as well as a keen interest in emerging trends within the machine learning space.

Contact us today to discuss.

Partner Group

Dwayne Price
Partner, Financial Services Advisory
T +353 (0)1 436 6494
E dwayne.price@ie.gt.com

Daniel Fernández
Partner, Financial Services Advisory
T +34 915 763 999
E daniel.fernandezd@es.gt.com

Amanda Ward
Partner, Financial Services Advisory
T +353 (0)1 433 2440
E amanda.ward@ie.gt.com

Nuala Crimmins
Partner, Financial Services Advisory
T +353 (0)1 433 2538
E nuala.crimmins@ie.gt.com

Brian O'Dwyer
Partner, Financial Services Advisory
T +353 (0)1 433 2538
E brian.odwyer@ie.gt.com

Frankie Cronin
Partner, Financial Services Advisory
T +353 (0)1 646 9044
E frankie.cronin@ie.gt.com

Quantitative Risk Model Development Group

Lukas Majer
Director, Quantitative Risk
T +353 (0)1 646 9006
E lukas.majer@ie.gt.com

Mark Perry
Director, Quantitative Risk
T +353 (0)1 408 6909
E mark.perry@ie.gt.com

Jonathan Fitzpatrick
Director, Quantitative Risk
T +353 (0)1 680 5805
E Jonathan.Fitzpatrick@ie.gt.com

Andreas Spyrides
Director, Quantitative Risk
T +357 (0)2 260 0270
E andreas.spyrides@cy.gt.com

Eleni Marinou
Manager, Quantitative Risk
T +353(0)1 680 5805
E Eleni.Marinou@ie.gt.com

Victoria Álvarez
Consultant, Quantitative Risk
T +34 915 763 999
E victoria.alvarez@es.gt.com

Alvaro Herráez
Consultant, Quantitative Risk
T +34 915 763 999
E alvaro.herraez@es.gt.com

Cormac Monaghan
PHD Student
Maynooth University

Toheeb Jimoh
PHD Student
University of Limerick

Offices in Dublin, Madrid and Nicosia.

grantthornton.ie
@GrantThorntonIE
Grant Thornton Ireland

Grant Thornton

grantthornton.ie