www.ijecs.in
International Journal Of Engineering And Computer Science
Volume 10 Issue 12 December 2021, Page No. 25671-25694
ISSN: 2319-7242 DOI: 10.18535/ijecs/v10i12.4676

Machine Learning Approaches to Credit Scoring and Portfolio Optimization in Wealth Management

Abhishek Dodda

Principal Product Manager, ORCID: 0009-0000-6728-945X

Abstract

There is a wealth of information available to wealth managers for analyzing clients' profiles and their financial needs. Machine learning techniques can help to extract better insights from this data. Using machine learning, wealth managers can make better predictions about clients' life events and financial preferences, resulting in better financial advice and improved revenues. Besides client profiling and advising, machine learning can also be used to optimize the asset management process. Machine learning techniques can now outperform other statistical models traditionally used in finance in both credit scoring and portfolio returns prediction, while also offering properties such as dimensionality reduction, transfer learning, and semi- or unsupervised learning that are not available in most competing models. Wealth managers can benefit from using machine learning to better predict default and portfolio returns, resulting in increased performance from better portfolio construction and rebalancing.

Although machine learning has received more and more attention in finance over the past years and has been successfully applied to a wide variety of areas, some other aspects have only recently started attracting attention. The majority of papers in the asset management field for instance still rely on classifiers traditionally used in finance, such as linear models. Moreover, even though they highlight the appealing properties of machine learning methods, there are still too many asset managers that remain skeptical about the usefulness of machine learning models and the risks surrounding their implementation. We explore these issues by applying and comparing several machine learning algorithms over real-world financial data for the problems of credit scoring and portfolio optimization. We show that most algorithms can indeed lead to better predictive performances, leading to well-founded financial decisions.

Keywords: AI, wealth management, machine learning, personalized investment, big data, finance, risk management, predictive modeling, cloud infrastructure, scalable services, agile finance, credit scoring, portfolio optimization, AI integration, cloud computing, wealth advisory, data-driven strategies.

1. Introduction

Wealth management is known to provide integrated, goal-based financial advice to individuals, trusted by them in the process of capital preservation and accumulation. The services offered by a private bank to help its clients manage their wealth typically fall in the area of portfolio construction, monitoring, and rebalancing. While a bank's clients typically rely heavily on the bank's financial advisor on portfolio construction and recommended asset allocations based on risk profiles, the advisor's recommendations are ultimately based on subjective human judgment, except for model portfolios based on general asset allocations for the risk profiles available. Building on some empirical studies in portfolio analysis, we argue that the recommendations could be improved further if the recommended portfolios were backed by systematic computations or predictions stemming from empirical studies. Further, since it would be silly to have the bank advisors recommend a single portfolio with no other portfolios for comparison for clients within the same risk profiles, it naturally leads us to the idea of portfolio optimization by different models and the associated, tailored advisory service by the banks.

In this paper, we present an overview of some of those portfolio optimization approaches with some machine learning flavor tailored towards current market landscape and conditions. Another track for financial advisors is to guide clients in asset selection. In the traditional optimization framework, the asset returns are typically modeled to be governed by some distributional assumptions, which does not hold true for actual markets. Consequently, most optimal portfolios in their approach suffer in terms of actual returns, and can lead to surprising results. These portfolios could be improved through predictive modeling of asset returns based on certain features representing current market landscape, as well through the use of more advanced methods for asset return forecasts.

2. Overview of Credit Scoring

Credit scoring is a significant tool for credit risk measuring and assessment in the financial industry and consists of a collection of processes that aims to evaluate borrowers' creditworthiness through an integrated use of statistical analysis and behavioral evaluation. As a pivotal component of the lending and credit discipline, it is beneficial for both lenders and borrowers. Implemented by lenders through a systematic and automated estimation process, credit scoring allows to distinguish deserving clients from those that will not pay back credit they owe, as well as to reduce the need of manual risk evaluation for each single applicant. On the other hand, by considering clients' specific profile as well as the available credit supply, increasing the scoring cutoff can favor borrowers who optimally match lenders' risk requirement. This reciprocal convenience has driven the growing utilization of credit scoring by both lenders and borrowers to facilitate lending and credit provision.

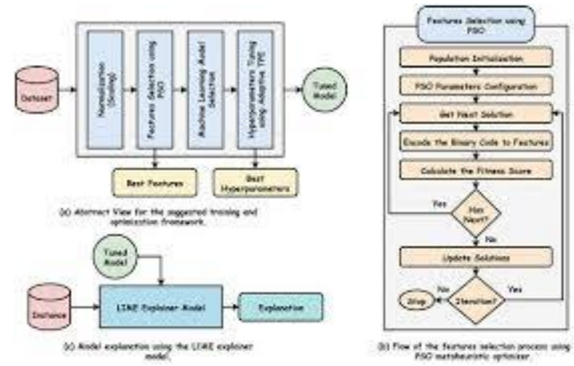

Fig : 1 Mathematical Modeling and Analysis

Credit scoring has rapidly spread in a variety of financial products, alongside the revolution of lending. Credit scoring was first adopted about 60 years ago, for home mortgage loans. Since then, it has gradually conquered the whole consumer credit sector, with a special mention for credit card loans. Today, credit scoring is present in more complex and innovative lending models, such as direct lending, peer to peer, and marketplace lending. More recently, with the advent of machine learning techniques to artificial intelligence, these models implemented through the usage of big data and methodologies have been greatly improved. Indeed, techniques applied to diverse platforms have allowed to obtain significant scoring results, also when they are applied on top of more conventional econometric scoring models.

3. Importance of Machine Learning in Finance

The increase in computational power, combined with the availability of immense amounts of data, has made machine learning a new tool for investors to maximize the profitability of their assets. Machine learning represents a powerful approach to find hidden structure in data sets, mimicking or surpassing human expertise. It is based on the idea that making accurate predictions about an unknown outcome can be improved incrementally with the use of data related to past observations. Machine learning can be considered a specialized subfield of statistics that studies and develops algorithms that can be trained on data to make accurate predictions. Fitting a model to data is often done in statistics, but machine learning has been developed to achieve different goals, such as exposure to a range of possible outcomes and automatic refitting when the properties of the outcome space change over time. In finance, many ML techniques have been applied successfully. Yet, it is worth noting that the use of ML in finance is not a solution to all the problems that financial practitioners face. Careful consideration needs to be given to the choice of the ML techniques. In recent years, the traditional statistical methods used to model financial markets have been challenged and augmented by machine learning models. These models have been shown to have superior prediction performance across various domains. As a result, practitioners are increasingly utilizing ML methodologies to automate and enhance various stages of the model development process.

4. Machine Learning Techniques

Recent developments in technology and computer sciences, especially in the fields of computer processing capacities, smart devices, and data generation and storage, have made available many sources of data. At the same time, we have seen the emergence of a new class of algorithms, which are grouped together as machine learning techniques. Machine learning, as an applied branch of artificial intelligence, aims to solve problems in domains of different nature, using similar algorithmic procedures. The domain of application is typically identified by the data used. These techniques have proven effective in many – if not most – decision and prediction problems in different sectors. More specifically, these same techniques are gaining traction since the turn of this century in finance, marketing, operations research, and economics. We next introduce a concise taxonomy of main machine learning techniques used throughout this work.

Machine learning techniques are typically divided into three groups: supervised learning; unsupervised learning; and reinforcement learning. Supervised learning applies to situations where one has a sequence of past records from which to learn. For each record, the label, or target, is known. The model is trained with these records and results, to make predictions on future occurrences for which the label is unknown. This process is called prediction. Most commonly, supervised learning is used when the value of the dependent variable is numeric. In this case, it can be thought as predicting the conditional expectation of the dependent variable. Predictions are generally obtained for all data points in the training set, validating the model against a sample not included in the model fitting.

4.1. Supervised Learning

The application of machine learning to problems where a response variable is present has sometimes been referred to as supervised learning. The decisive aspect is that the algorithm that learns the relationship between the explanatory variables and the response variable, often referred to as the target variable, does so using many examples where the outcome is known. This is typically referred to as training the algorithm, which can then be applied to examples for which the outcome is unknown. Supervised learning is therefore a mapping problem, with user-specified rules governing the quality of the achieved mapping. These need only apply to the training dataset, and methods often operate internally to avoid overfitting, although out-of-sample testing is normally de rigor. Supervised learning provides a predicted label for instances only falling into classes or types corresponding to those in the training data, whereas in generalization methods these classes or types can be generated by varying the algorithm input values.

Eqn.1: Logistic Regression (Binary Classifier for Credit Risk)

$$P(y=1 \mid \mathbf{x}) = \frac{1}{1+e^{-(\mathbf{w}^\top \mathbf{x}+b)}}$$

- $\mathbf{x}$ = feature vector (e.g., income, credit history, etc.)
- $\mathbf{w}, b$ = model weights and bias
- Decision rule: Predict default if $P(y=1 \mid \mathbf{x}) > \tau$

An important differentiating aspect of supervised learning is the need to quantitatively measure the quality of the variable mapping. This is necessary not only in order to define the goal of the learning or training process, but also to carry out model selection of the many candidate algorithms, often from very different classes, and to calibrate the selected algorithm. For supervised learning applications, many such predictive measures are available, defined as the expected value of a function of the input data and the output. In general, the function evaluates the predictive accuracy of the algorithm when applied to test data falling in a region of input space, weighted by the probability distribution of the input space region. By focusing on parameterized functions, known as loss or scoring functions, it is possible to define many model evaluation metrics, based on assumptions on the probability distribution. These include classification accuracy even for multiclass models.

4.2. Unsupervised Learning

This section summarizes key methods and techniques in machine learning that do not belong neatly in the previous section on supervised learning and their application to investment issues. As we point out in the introduction to this chapter, machine learning methods are often presented using very specific problem settings. However, while such presentations are informative and enable researchers to implement their methods into problems, simply because of the simple and low dimensional nature of many of these examples, ideally, we would like to make this presentation more systematic and thereby allow an implementation of these various techniques across many different disciplines out of the box so to speak. However, this is a long-term goal of our project and would require substantial more detail than is currently possible with the work of only a few individuals.

The techniques thus far summarized are a small subset of the terminology and variety of tools available in machine learning today. However, they have been found to work on many real-life problem settings and we would like to offer them in a consistent manner to researchers and practitioners alike so that they are able to transfer algorithms across different disciplines and issues. Especially in the investment area there exists a large portfolio of techniques to choose from. Outside of the wealth management framework, many of these techniques have been used already by researchers to study investment issues. But mainly using only a small number of the tools available and thereby also limiting the findings researchers are able to achieve. Hence, this is a small contribution to the financial engineering that we foresee within the quantitative finance industry using the tools of machine learning.

4.3. Reinforcement Learning

Reinforcement learning is a type of machine learning that interacts with the environment not to infer a function from data but to learn from the consequences of its actions on the environment. The learning agent is in a state at time t and takes an action A(t), changing the environment state to S(t + 1), and receiving, based on this transition, a reward R(t + 1). The action is selected according to a policy or behavior function that provides a mapping from observed states to action probabilities. This mapping is in general stochastic, and the control objective may be to optimize the value of discounted rewards over the long term. The problem is formalized through a Markov decision process that describes the environment dynamics by determining the probability distribution of future states, given the current state and action. Value-based solutions learn a value function that provides a performance measure over time for each action when applied from a state. The Q function gives the 1-step value of executing action A in state S, and the objective is to find the policy that maximizes the expected value.

Reinforcement learning is typically used for problems where supervised learning methods are not sufficient, like dynamic pricing, options market making, inventory control, or trading strategies. In these applications, traditional control approaches are often not able to exploit the real data, which are needed to solve the problem. The flexibility offered by the reinforcement learning combination of models with data allows applications to more complex systems, as the reinforcement learning solution becomes a hybrid of empirical and theoretical modeling.

5. Data Sources for Credit Scoring

This chapter discusses datasets underlying traditional data sources for credit scoring and reviews some novel datasets labeled as alternative data sources. A credit scoring model that aims to predict an event of interest must be developed based on a labeled dataset where the observed features relate to the predictive features and the labeling model specifies the information about the event. Traditional datasets are credit registries provided by national banks, credit bureaux and other organizations providing credit information. These registries provide detailed information about the applicant's previous credit history and payment delinquencies, links with other borrowers, and demographic information about all considered applicants. At the same time, data for a very limited scope is provided. Therefore, designing a predictive model trained only on traditional data can yield inconclusive results, especially if the applicant is a new customer of the credit institution. Alternative data sources catalogue different groups of data sources labeled as alternative data sources. Altogether, alternative data sources aggregate information about an applicant's behavior across many important aspects beyond traditional data. Information from alternative data sources can be used to extend traditional registries or can be used solely. Extending traditional data can positively change labeling of the event of interest in the future, thereby enhancing model performance. There is still a need for reliable country-specific credit scoring datasets.

5.1. Traditional Data Sources

The key to a correct credit decision is an appropriate estimation of the default probability of the borrower. In the conventional credit scoring scenario, one would estimate a scorecard by using past borrower performance and a set of particular characteristics of borrowers such as credit history and behavior data, age of customers, type of residence, employment status and history, credit bureau scores, types of accounts, loan to value ratios, etc. These characteristics can be extracted from three traditional sources of data, namely loan performance data, credit bureaus, and behavioral data.

One may analyze the loan performance data of a particular institution to obtain particular lending insight. The performance data indicate how and when in the past the institutions' borrowers defaulted on their loans. By merging the performance data with demographic information about the borrowers and other internally available characteristics, one can use logistic regression or other statistical techniques to estimate a scorecard. However, institutions do not always have access to a comprehensive collection of past borrowers. Credit scoring based on such cohorts may inadvertently exclude some of the discriminating variables needed for predictive purposes. Furthermore, such a scorecard tends to focus on pre-existing customers, rather than new applicants. Such cohorts may also be too small to lend sufficient credibility to the scorecard estimation. The completion of the data collection will generally depend on the classification of loans. Credit performance data is often inexpensive and quick to access. Implementation of a scorecard based on internal data is straightforward.

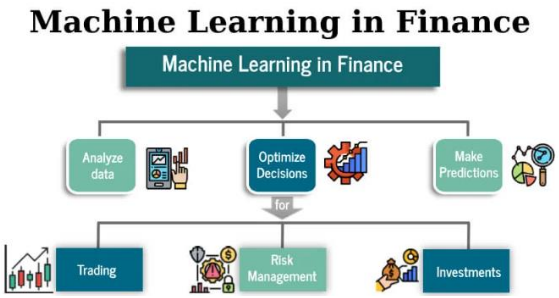

Fig : 2 Machine Learning in Finance

5.2. Alternative Data Sources

Finding information on an individual or business is not a straightforward endeavor and has been traditionally limited to official sources, such as credit bureaus. In recent years, new data aggregators, thanks to the advance in computing and data storage technologies, offer a variety of data sources and a more detailed profile for customers, businesses, or industries. In fact, the growth of interest in alternative data sources used for credit scoring is linked to several ongoing qualitative changes: the augmentation of public interest in financial inclusion; the major innovation of FinTech companies, which provide online or mobile credit applications and services; the negative impact of the financial crisis, which prompted banks to favor the repair of their balance sheets rather than those of their clients; the pressure imposed by the competent authorities to prevent the eruption of systemic crises caused by the defaults of banking clients. The idea is to reduce information asymmetries by using these alternative data sources to determine risk scores in order to allow banks to lend to those inside the credit report "dead zone."

Many lenders use only credit files that depend on credit bureau activity, while others utilize detailed income verification, automated underwriting systems, and loan application data, as well as tradeline databases. Most commercial banks also trust directly to traditional data sources, using credit bureau information like that used to feed scoring tools for decades. Nevertheless, a number of nonbank financial institutions have at disposal both traditional and alternative data to make their lending decisions. With respect to those who use alternative data sources, some incorporate it in their traditional models, others used advanced machine learning methods that make it possible to extract information in a more complete way, while some others create new models specifically for alternative data.

6. Feature Engineering in Credit Scoring

Credit scoring is one of the most fundamental applications of machine learning in risk management. Banks use credit scoring to automatically and objectively assess customer creditworthiness as well as predict losses from loan defaults. Credit scoring ensures that banks follow the principles of proper risk assessment while following regulations and preventing discrimination. Automatic and objective credit assessment leads to an efficient allocation of the entire credit supply to customers while preventing discrimination of customers. As the customer base is often heterogeneous, segmentation of the credit demand is important to apply differentiated and risk-adjusted pricing. Automating and standardizing the credit assessment process enables the linking of the risk and the pricing of the credit.

An important part of the credit assessment process is determining the influence of the single features as well as parameterized or non-parameterized feature transformations on the risk probability but also the probability of a loan default. Typically, these transformations are determined with limited granularity and without learning. Through machine learning, explanations or expert rules can be automatically derived, thus discovering interactions and relationships which are otherwise overlooked, thereby optimizing the modeling of the credit scoring process. These findings from machine learning will have a significant impact on credit scoring as well as the decisions taken by banks in the future. Improving these ads and also the decision process will enable the banking leads to enhance and improve the quality of customer relationships and thus help to increase the willingness to pay for banks.

Decision trees, random forest, boosting as well as neural networks enable the automation of the decision process used in risk management. Modeling the decision attributes correctly is paramount for sound and optimally performing loan default prediction processes. Although many of the machine learning approaches are hailed for their automatic feature creation or transformation processes, it should be noted that this statement is somewhat exaggerated. While many of the machine learning methods do allow the use of transformed or improvised features, extensive feature engineering processes are the cornerstone of sound decision model emerging from any of these machine learning approaches.

6.1. Feature Selection

In the literature of potential predictors of credit risk default, the number of proposed features is enormous, as it includes information not only from financial features, such as bank statements or credit scores, but also population features, such as demographic data or even affinity features, referring to data from the social and activity networks to which a client belongs. Feature selection from this massive pool of candidate features is thus a critical task for credit risk modeling.

Three important observations in this setting are discussed. First, data scientists usually have a significant freedom of choice within this pool of features. However, prior knowledge about the relationship of some data with other related fields can definitely help to choose more effective features for credit scoring. Second, feature selection is not easy. Dimensions may vary in a highly non-linear fashion, there might be many irrelevant features, multiple groups of correlated features may exist, features may affect different data points in different ways. Adding features can decrease the amount of noise in the data but may also make the task of constructing a credit score more complicated, leading to larger scorecard and misclassification of unseen data, and potentially lower predictive accuracy during operational use. Predictive accuracy may be boosted on the other hand by adding univariate transformations, such as interaction terms or dummy indicators of categorical joint distributions. Moreso, it could be the case that a less precise model is requested, with a higher speed of execution for decisions. This could be the case for exploratory assessments of risk probabilities in the portfolio analytics phase. When a model is used at decision time at the financing moment, accuracy is indeed critical.

6.2. Feature Transformation

We focus on how to optimally represent the chosen features for the downstream machine learning model. Feature transformations are used to change the feature into a different scale or to create new features from the existing ones. Feature transformations are often cheap and help machine learning and non-machine learning methods alike. Therefore, we stress feature transformations since they can considerably improve performance. Furthermore, we augment existing features so that the estimator learns a richer structure yet trains fast on the raw data. Finally, we investigate how to optimally encode categorical variables and other surrogate feature transforms.

Bagging and boosting algorithms have the potential to introduce non-linearity structure if needed. We survey some existing transformations, but none is exhaustive. In practice, we often need to rely on intuition, domain knowledge, reviewers, and engineer better features. Credit scoring features are usually engineered to be well-behaved, that is, result in reasonable decision-making rules. Arbitrary generalization may temporarily outperform such base features, but it is unlikely to transfer well to real decision making. Candidate Transformation Method - Create new ordinal variables from real-valued variables – Knapsack Variables, Scores - Create new real-valued variables from ordinal variables – One-Hot Encoding - Create Products or Interactions of Variables, Generalized Additive Models - Fusion, Smoothing Operations, Coarse Binning via Clustering, Piecewise Constant Models such as CART – Composite Score, Non-Additive Functions, for example: Negative Slope Loss Function - Create Indicator Variables for Derived Variables such as Defaulter Age – Adding Age-Interaction Variables, Bureau Mix, Credit-Payment Mix, and Other Derived Credit Variables - Reserve Borrowing Limits for All Borrowers – Reserve Upper Borrowing Limits for Special Borrowing Groups.

7. Modeling Techniques for Credit Scoring

We explore and discuss four standard modeling techniques considered for credit scoring, in the order of ascending complexity. Logistic regression is a basic parametric technique that has been commonly used since the inception of statistical credit scoring and continues to be popular due to ease of interpretation. Decision trees are the next step up in terms of applicability and sophistication. Random forests is a generalization of decision trees that is especially potent in its predictive potential. The final model to be explored is an artificial neural network, which dates back to the 1950s, but has only recently been enhanced to be deeply layered and more powerful. Neural networks are particularly well-suited for image, sound, and text-related pattern matching, but in the modern era of big data and advances in computing power, have found much wider applications.

Logistic regression remains a popular tool of statistical analysis, especially when results need to be interpretable by non-technical users. In its basic form, the logistic regression function has the following elegant and interesting mathematical form. The logit model can also be estimated in the context of multi-logit probabilities, where the user has more than one option among several possible choices. The fact that our probabilities sum to one then places a linear constraint on the many logit "betas" involved, thereby allowing for estimation of such models using various familiar econometric methods.

7.1. Logistic Regression

Logistic regression is one of the oldest models to classify the labeled data. Named for the logistic function used to convert continuous predictions to class probabilities, logistic regression is able to predict the probability of categorical dependent variable with $k > 2$ classes or two classes for the case of binary dependent variable. In this basic model design, the dependent variable is binary while the independent variables can be either continuous or categorical variables. The logistic regression model is usually estimated by using maximum likelihood estimation. Logistic regression performs well in terms of classification accuracy while requiring lower computational resources compared to other approaches. It is widely adopted despite availability of more sophisticated classification algorithms.

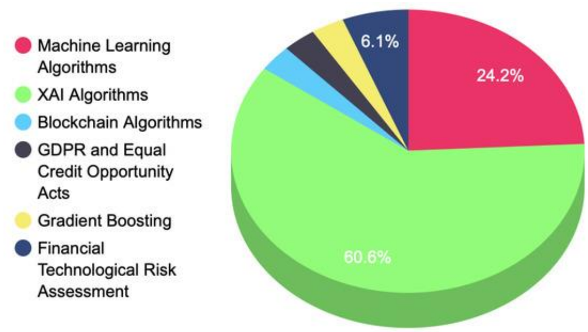

Fig : Support Using Explainable Artificial Intelligence

However, unlike other predictive models, logistic regression does not provide non-linear boundary fit and requires large sample size. It is important to ensure that the assumptions of the logistic regression model are satisfied. It assumes independence of variables, absence of multicollinearity, no outliers either in dependent or independent variables as well as linearity between logit of independent variables as well as dependent variables. The relationship between independent and dependent variables is additive and multiplicative. Logistic regression allows straightforward estimation of class probability and works well if there is a logit of class probabilities that is linear in the predictor variables. Additionally, the model can work with outliers and performs better than the more complex algorithms with small sample sizes. All of these properties make logistic regression popular for classification in some cases. Its simplicity and transparency make it an excellent first modeling approach in credit scoring.

7.2. Decision Trees

Decision Trees (DT) are a supervised ML technique that derives its name from the structure of the algorithm. Decision Trees work by continuously splitting the data into subsets based on the value of each feature. The splits are done in an informed manner, meaning that the algorithm chooses the attribute splits that perfectly and optimally separate the data. The algorithm chooses the best way to branch the data based on the Minimum Description Length principle, also known as Kolmogorov complexity, based on Shannon's entropy function and information gain measure. Using these measures, the Decision Tree algorithm identifies splits that partition the samples into purer intrinsic classes. The process continues recursively until a stopping criterion is reached, resulting in an inverted branching structure with multiple leaf nodes that correspond to the different classes. If the leaf nodes correspond to the target sample, then it is known as a perfect tree.

Eqn.2: Decision Tree Splitting Criterion

$$G(t) = \sum_{i=1}^{K} p_i(1 - p_i)$$

- p_i is the proportion of class i in the node

The tree construction is a greedy search that begins at the root node and branches the data iteratively into the child nodes. The algorithm chooses the attribute to branch upon into child nodes based on the impurity measure considering all the available features. After forming each child node, the algorithm continues the process recursively until it reaches a terminal node, which signifies purity or a stopping criterion. Decision trees can be built cementing different characteristics, primarily the splitting criterion, the tree depth, the number of available features to select at each stage, and the stopping criterion. The most common impurity measures are information gain, Gini impurity, and misclassification error. The combination of these parameters generates different decision tree methods.

7.3. Random Forests

In this section, we focus on the Random Forest, a popular method that combines functions of decision trees to achieve better generalization performance. Random Forest is an ensemble that combines the predictions from multiple base learners. The inputs to the Random Forest are parameterized weak learners, where the idea is to use the voting mechanism to aggregate the predictions.

Let P be the maximum number of classes for a classification task and H be the output probabilities of a new sample. Then, when we take the classification with the maximum confidence, the expected classification error is expressed as the sum of the error with the voting mechanism, the estimation functions using histogram and the estimation error of the multi-class probability using bounded variation functions:

Both the second and the last terms will converge to zero as the number of sampled base decisions gets larger. Besides, the histogram estimation will capture most of the functions in the estimation space. Thus, it is enough to take an appropriate number of weak classifiers to reach our approximation. Note that if we take these to be decision stumps, then we also have a bound for the number of decision stumps required for the approximation.

7.4. Neural Networks

Neural networks are one of the most elaborated machine learning and artificial intelligence methods. Their theoretical foundation go back fifty years ago, but the first practical applications appeared only when enormous data and investigation resources were reserved to specialized teams ten years ago. Commercially available software packages have improved tremendously since, and neural networks can today be applied rather easily without elaborate mathematical preparation. With neural networks, the general tasks of modeling and forecasting are approached with various architectures of feedforward or recurrent networks. They are connected units that mimic the information processing of biological neural networks. Such networks have in common a flow of information in one direction. There are several hidden layers of geometrical configurable neuron units processing in parallel the weighted inputs from the previous layer and passing on to the next layer the results of a non-linear activation function. Technically, neural networks are universal function approximators that require no distributional assumptions. Neural networks can be trained, after presenting them many examples of the mapping between inputs and expected outputs, by an optimization with gradient descent. Although there are different algorithms to update the weights, the most widely used is the backpropagation algorithm, in which the weight updates are determined by the partial derivatives of a cost function with respect to the weights. This approach is far from optimal, given that the gradient does not exploit the fact that the individual terms in the cost function are not independent. Instead, it uses second derivative information from all examples. Hence, it is recommended to use mini-batch algorithms that combine aspects of the stochastic and full-batch algorithms. Neural networks are popular because their prediction accuracy is among the best, and their versatility allows one to complete several tasks, as detection, classification, regression, transformation and dynamic modeling.

8. Evaluation Metrics for Credit Scoring Models

In order to assess the quality of credit scoring models, one has to define metrics to compare the predictions with the real outcome of loan status. In a classification problem with a binary outcome, as is the case in credit scoring, the best prediction is obviously a model that predicts perfectly for both classes. However, although it is possible to create models that classify every borrower correctly, this is often not feasible and fails to consider the goal of the task, which is to differentiate between good and bad credit risks, and the costs involved in doing so. A very simple metric is the total accuracy, which quantifies the fraction of loans correctly predicted by the model. However, in the domain of credit scoring, labels are often highly imbalanced. As an example, one can consider historical data where a certain percentage of customers default and another percentage of customers repay their loans. In this case, if one percentage were 0.5% and the other were 99.5%, a naive model that predicts all customers to be of the second class would achieve an accuracy of 99.5%, while doing a poor job of predicting bad credit risks. In this case, using accuracy as a metric would not only lead to a model that is of no use in practice, but also potentially to a model that performs worse than merely classifying all instances as the second class and not learning anything from the training data.

A very common metric used in practice to measure model quality is the confusion matrix, which distinguishes between the true positive, true negative, false positive, and false negative predictions. Based on this, one can compute the model precision as true positive divided by the sum of true positive and false positive, the recall as true positive divided by the sum of true positive and false negative, and the F1 score as 2 multiplied by (precision multiplied by recall) divided by the sum of precision and recall. Precision is defined as the number of true positive predictions divided by the total number of predictions made for the positive class. In our case, that characterizes how many of the predictions for the default class are in fact components of the default class. Therefore, this is a good metric to apply when defaulting is costly and one wants to make sure that a loan predicted to default does indeed actually default, which would usually be the case if one employed such a model in a real lending case.

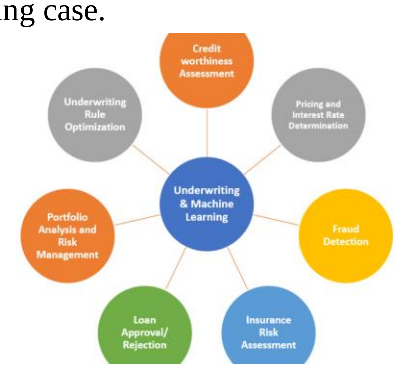

Fig : 3 Underwriting & Machine Learning

8.1. Accuracy

Model evaluation is an essential part of any modeling task in machine learning and artificial intelligence. It is also a critical part of fraud detection models. If a model is not robust and performs poorly when it is put into production, the organization will lose money and time rectifying issues. Fraud detection issues have class imbalance, so unique metrics are paramount to evaluate whether the models will perform robustly with the imbalanced data. There is also a risk that the evaluation will be contingent on the exemplary data used, which may introduce a potential bias. The performance metrics then play an essential role in linking the model to the business's underlying goals.

Accuracy is a model evaluation metric that measures how many predictions did a model get correct. Considering a two-class problem (0,1), it is defined as

ACCURACY = (TP + TN) / (TP + TN + FP + FN)

However, this metric in a fraud detection setting would be misleading as fraud detection tasks typically have a high imbalance between classes. For example, let's consider a task in which we want to detect a rare disease that affects only 1 in 100,000 patients. If we simply predicted all patients as sick, we would achieve an accuracy metric of 99.999%. However, this does not reflect the underlying truth of the predictions: these would all be false positives, and no one would actually be sick. Based on this consideration, we can see how accuracy is not a reliable metric to employ when we have imbalance in the classes. It does not provide additional information regarding the classes' correct predictions other than that the total prediction count was right.

8.2. Precision and Recall

Precision is defined as a ratio of true positive observations to total positive predictions. Recall is calculated as a ratio of true positive observations to total actual positives. Both precision and recall provide information about the false positives and false negatives rates, which assists in understanding the type of error present in the prediction problem. This analysis is particularly useful when working with a highly imbalanced dataset.

Applying a credit scoring model on a large pool of clients can easily yield a skewed distribution of predicted loan defaults versus the actual loan defaults. In such a case, a large portion of the clients would get predicted to not default on the loan. Therefore, a naive classification model yielding such accuracy will have a high predicted to actual ratio and consequently high accuracy. However, precision and recall will help refine the process by checking the ratio of predicted defaults to actual defaults. Such an analysis is crucial when the defaulting or terminating event is an extremely-sensitive event such as an account defaulting and the cost of prediction error carries a significant recrimination cost.

8.3. ROC-AUC

An intuitive way to evaluate the performance of a classifier is to fix its decision threshold at each possible value on its range in the interval [0, 1], which is normally uniform and equally distributed, and, accordingly, compute the accuracy, precision, and recall, where TN and FP are true negativity and false positivity, respectively. A curve showing the drop in accuracy, precision, or recall, as a function of the threshold value would be generated. Precisions@t and recalls@t will drop when the target population becomes large because the remaining majority formed of nontarget will inflate the TN value, while acc@t will drop when the count of remaining FP becomes large because the underlying population will be equally weighted since TN and FP will be tightly linked. Using standard additive convenience selected with a desired t-bias will return convenient utility functions corresponding to having the maximum positive payoff or minimum negative loss, controlling both ease and risk, or the maximum gross yield using precision@t, or the maximum net ease using recall@t, or the tie between both using acc@t. Posterior probabilities will be generated according to whether the observation Yi = 1, or Yi = 0.

When an input observation is predicted as not belonging to a target class, we are not too concerned about the actual classification. This is because scoring methods are, somewhat, a finer type of classification methods, and in most applications, the predicted probabilities of a method are more valuable than the predicted classes. The ROC curve plots the tradeoff between sensitivity and fall-out (or “false positive rate”) on different threshold levels, which can be summarized succinctly using AUC. Thus, by using the machine learning probabilistic outputs, we have an effective weighting scheme to compute sensitivity and fall-out, and the same fluctuations occur in space stemming from the underlying asymmetrical class distribution and the variety of probability distributions associated with the two classes.

9. Challenges in Credit Scoring

Machine learning can help banks, lenders, and other financial institutions leverage the big data revolution, assisting with customer selection, screening, risk assessment, risk mitigation, product pricing, early warning signals for default, and related tasks. Credit modeling with traditional statistical methods, however, has been continuously criticized for overly relying on feature engineering, lack of data transparency, rigidity, and reduced flexibility. Are the purported advantages of machine learning indeed true in these applications? Unfortunately, it is found that there are very substantial challenges in machine learning, specifically in data quality and accountability in its credit scoring applications.

Doing credit modeling well demands good, clean, and reliable data on credit clients. Having been in their businesses for substantial time periods, banks and other financial institutions may show their experiences in gleaning good and reliable data sets. The increasing popularity of the digital economy certainly offers opportunities for information to flow into big data lakes, enriching credit scoring databases. Yet it is abundantly clear that credit modeling without dealing with poor data quality will not improve the model performance, no matter how impressive the algorithms. In fact, the effect of poor data quality on model development can be even more magnifying in the case of machine learning based methods because of their inherent reliance on data and emphasis on automatic learning. So, data quality is a critical issue that credit model builders should pay sufficient attention to. With the considerable reliance on data in machine learning applications, there are serious concerns of model bias and fairness.

9.1. Data Quality Issues

The advent of Big Data technology arrived with the optimism that our capability of collecting information about consumers would make it easier to create accurate predictive models. However, paradoxically, we now find ourselves working within an environment where data quality issues are affecting our ability to build reliable Artificial Intelligence-based predictive models. Many machine learning algorithms can throught to be "magic boxes", which are capable of learning hidden structures in data without human intervention and therefore without considering the validity and information content of input variables or the general reliability of input data. Hidden biases in input variables, as well as noise or missing values among the data, can severely affect the quality of resulting credit scoring solutions and portfolio predictive models. These effects are then transferred to the credit risk management policies targeted by artificial intelligence.

The issue of missing values in the financial domain is often addressed with the assumption that they are missing at random and often with a random forest imputation. A more serious issue is the presence of mislabeled variables both in samples that are already available and that are used as input for the learning algorithm or even more seriously in the labels that determine the target population in the models that are then evaluated with scorecard methods. The traditional solutions for these issues have been to design monitoring models aimed at verifying the stability over time of estimated models and their predictions, validating the decisions taken by the scores assigned through the real default experiences of the borrowers. These traditional solutions can still be considered a tenet in this domain. However, new robustness checks have also been proposed that can be implemented to address the several issues mentioned above and help in the creation of interpretable predictive solutions in the credit domain for portfolio credit risk applications.

9.2. Bias and Fairness

Bias against disadvantaged or vulnerable groups is a long-standing issue in financial services and lending. Some practices may reinforce an existing bias, while others may instigate one. Historical lending data are often biased toward a particular demographic or minority group that has been discriminated against in the past. It is a just concern that specific features may request higher leverage from certain population subgroups, which may indeed induce bias discovery. It is, therefore, important to understand how bias measures work in general and how fairness is handled in credit scoring, which is especially sensitive to the questions of fairness. Suppose a data-driven credit scoring approach exhibits any sort of bias against a particular demographic group. In that case, financial institutions are under pressure to rectify the outcome with guidance from policymakers.

The goal of establishing a sectoral regime is to ensure fair and equal access to financial services for all individuals regardless of race or ethnicity. Financial services are also characterized by large amounts of personal information available to financial institutions. By having access to this selection of data, bias can occur via two different dimensions: model selection criteria, and input data selection. Selection of input data can be executed according to either quantitative or qualitative criteria. Selecting data with biased features would result in the learning of biased models reducing access for disadvantaged groups. The opposite is also true: on the one hand, ethnic bias and gender bias can occur both independently and jointly in a prediction or classification model for creditworthiness. Ethnic balancing for attribution and gender balancing for balance prediction seem to be important fairness measures. The joint objective of both credit modeling systems is to produce comparable scores and classification probabilities for both ethnic and gender groups. The inputs to the respective models differ only on a few predefined bias-related parameters.

10. Portfolio Optimization Techniques

This section describes the portfolio optimization mechanisms, which enable determining an optimal structure of the total wealth of a client. The goal is to determine an asset mix for which the expected portfolio return is maximized, given a threshold mortgage risk. Alternatively, the desired expected portfolio return is determined, such that the selected level of risk can be achieved. A risk/return might be implicitly or explicitly defined. In the base case, a quadratic utility function is specified, which leads to a mean-variance formulation. In a more general case to a higher moment function.

10.1. Mean-Variance Optimization

The original framework was developed in a series of publications based on earlier work, and it is commonly referred to as mean-variance optimization. The fundamental idea behind this approach is to maximize the return of a portfolio while being constrained by a taxonomic risk. For constructing a research fund through capital-for-a-city finding appearance with a minimum explication excess correlation, two budget constraints are considered. In this worry, proposed portfolios are minimizing the correlation with the main strategy. With equal portfolio weights, without volatilities and correlation, generally the proposed city portfolios are not risk-adjusted, giving that the mean deviation optimally sought for the active portfolios at a first step.

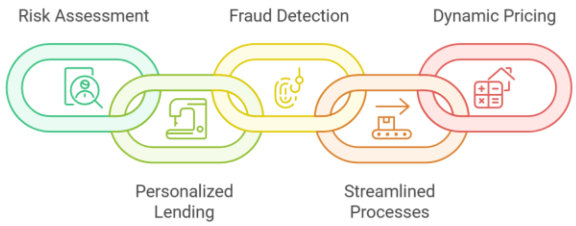

Fig : 4 AI-Based Credit Scoring

10.2. Black-Litterman Model

Despite its disadvantages, mean-variance optimization is still regularly used in practice, due to its ease of communication and implementation. Alternative asset allocation frameworks have been discussed, developed, and analyzed in the academic literature, the Black-Litterman model also falls into this category. The Black-Litterman model combines the original mean-variance optimization with the reverse optimization procedure. As such, a baseline patient reaction function for the return is validated with the implicit mean-variance risk aversion patience using existing portfolios returns.

10.1. Mean-Variance Optimization

Optimal asset allocation is a key component of a sound, well-performing investment strategy. Mean-variance optimization (MVO) is a widely recognized analytical framework for risk-return portfolio allocation and has been the foundation of portfolio optimization for well over 70 years. MVO assumes that returns on risky assets follow a normal distribution, that market participants are rational with homogeneous expectations about asset returns, and that they allocate their wealth to minimize the portfolio variance for a specific level of expected return. Therefore, MVO identifies the subset of feasible portfolios that minimizes the variance of the portfolio return for a given level of expected return, forming the minimum variance frontier, and the optimal set of feasible portfolios that maximize the expected excess return for a given level of risk, forming the capital market line. These two sets of portfolios are connected to the global minimum variance portfolio, which corresponds to the most risk-averse investor, located at the left quintile of the mean-variance frontier.

The optimal portfolio allocation is then determined with respect to investors' preferences for risk and return. Portfolio optimization in this context is generally aimed at the identification of the optimal risky asset allocation. Taking the investment horizon of the investor into account, an asset will be risk-free in the sense of generating a known return in the short term. Such a risky asset allocation means that every investor, regardless of her risk aversion or utility function, should have the same allocation across all currently efficient asset classes, which behave in a similar manner within the investment horizon of the investor. Whereas MVO provides a theoretically sound and intuitively appealing basis for constructing portfolios, its practical application in portfolio management has raised serious concerns. In particular, portfolios optimal for hypothetical markets are often poorly diversified, concentrated in a few risky assets or asset classes, and excessively invested in relatively risky assets.

10.2. Black-Litterman Model

Black and Litterman built on the mean-variance model of Markowitz to take into account the views of investors in addition to previous point estimates of the expected returns of assets and their covariance matrix. With respect to the mean-variance optimization model, the returns are more informative, as the model suggests that the investor allocate most of his wealth into one asset class as long as there is at least one asset whose expected return is sufficiently high. An extension of the Black-Litterman model proposed a more general model that does not require that the assets satisfy the Capital Asset Pricing Model, with which the original Black Litterman model is supposed to be compatible; and allowing estimates of the demand for commodities to be negative. The return forecasts that should be plugged into the portfolio optimization model so as to reflect the opinions are built by modifying the opinions by the investor's degrees of credibility and whether the position is directed in the

same way as the opinion is classified into opposing the expected excess returns or reinforcing it.
This Black-Litterman filter is a fascinating process to implement in practice, but there are few available pieces of software that can handle it under the extension proposed. In this section, we are going to implement the Black-Litterman model. In each step, we will also generate the preceding step inputs. In a first instance, the supervised model will behave as in Markowitz optimization. Indeed, we first assume that the return forecasts are only plugged into the portfolio optimization model, as if no opinion were related to them.

11. Machine Learning in Portfolio Management

11.1. Predictive Analytics

Investment predictions have the potential to facilitate more informed decision-making; behavioral finance points to human biases as detrimental for rational investment decisions. Once trained on historical data not affected by event risk, econometric models can be employed as a first step toward mathematical finance objectives: for example, the prediction of security returns based on their relationship with state factors can constitute a foundation for mean-variance or risk parity portfolio construction. However, the established econometric literature has delivered hardly predictable models for asset returns. By transferring and enhancing the recent predictive analytics success in different fields on portfolio analytics, ML has the potential to improve upon classic econometric specifications. In finance, recent applied modeling work along both the equity and the fixed-income dimension has leveraged on the hidden step-wise structure present in asset returns.
Indeed, to expect substantial returns above transaction costs from passive or naïve investment strategies based on predictive modeling risk-sharing systems needs permanent applicability, and thus returns have to be seen as investment premiums. Certainly, it seems prima facie reasonable to expect the existence of predictive investment signals or rules across short time horizons. Short time horizons are precisely the ones that most closely resemble digital signals, rather than technical indicators tracing psychological or behavioral trends. More involved ML setups could today be incorporated in portfolio management by financial institutions delivering AUM services for demanding customers with high investable wealth.

Eqn.3: Gradient Boosting Objective

$$\mathcal{L} = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k)$$

- l: loss function (e.g., logistic loss)
- $\Omega(f_k) = \gamma T + \frac{1}{2}\lambda \|w\|^2$: regularization on trees

Portfolio management has to define a sufficiently differentiated investment policy objective at one or more allocation decision horizons by taking into account the client-specific preferences and risk aversion profiles, as well as the actual risk-return distribution of the portfolio, defined by the relative size of the tail regions. Decision-making has to be followed up by periods of underlying portfolio structure adjustment triggered by transaction costs considerations for exceeding drawdown thresholds, for volatility breakouts, or for horizon-dependent conditional covariance levels with respect to risky assets, safe havens, or finally, other portfolios.

11.1. Predictive Analytics

Predictive analytics refers to a set of statistical techniques to analyze current and historical facts to make predictions regarding future events. Since classical statistics focus mostly on estimating parameter values, explanations, and cause-effect reasoning, statistical data mining plays the key role in predictive analytics. Instead of estimating relationships between numeric variables, or even the trajectories of metric-valued curves, predictive analytics aims to specify, and ultimately extract, the hidden information in datasets, making more accurate predictions possible.
Predictive analytics encompasses a wide range of predictive techniques: regression and time series models, multivariate adaptive regression splines,

generalized additive models, neural networks, support vector machines, decision trees and ensembles experts, nearest-neighbor techniques, and many more. As these techniques are diverse, they are selected according to critical analysis, experience, and availability of software packages. Note that, even if different predictive tools are based on different methodologies, they often provide similar predictions. Robustness is a highly desirable property in predictive analytics because the future is uncertain and the data are likely to be noisy. It is increasingly accepted that combining multiple diverse predictive models, known as model averaging or stacking, often yields better predictions than any chosen model. Predictive analytics is one of the main orchestras in the risk-return tango since the required constituent risk-return performance functions are typically unavailable. Most of the time, parametric specifications are too restrictive and do not fit properly the observed empirical evidence.

11.2. Risk Assessment

Understanding and quantifying the risk of an investment portfolio is essential for the formulation of an investment strategy. Any investment strategy introduces additional risk relative to a benchmark, be it positive or negative. The risk return ratio of portfolios is crucial to investors when deciding on the acceptance of additional risk. Subsequently, the importance of measuring risk is emphasized for each investment strategy. Predicting portfolio risk is claimed as one of the most crucial tasks within quantitative portfolio management. Probabilistic risk forecasts and extreme risk forecasts are often used by practitioners for risk calculations.

A general task for quantifying risk within portfolio management is to predict tomorrow's returns for multiple assets. However, for risk modeling this situation has to be treated differently from the return predicting case. Whereas portfolio returns are used to specify and evaluate the performance of a given asset allocation, it is the distribution of returns that drives the risk budgeting and risk controlling approach of modern portfolio theory. Probabilistic return prediction with a multi-label architecture is found to be superior to an independent single-label forecast strategy. A further approach to modeling multivariate return distributions from the perspective of scale location scale mixtures of normals is offered using multivariate projections. Probabilistic multivariate forecasts of a time-varying multivariate normal is found for the embedding of multiscale time-varying correlation structure within forecasted using a deterministic two-stage approach.

Predicting risk probabilities policies induce is another important contribution of supervised machine learning. We apply supervised neural networks to predicting the expected risk probabilities. Probabilities of returns exceeding specified percentiles are predicted simultaneously based on four predictors rank psychical integrated average scores. It is shown that the predicted risk probabilities induce policies that are competitive with the best performing strategies.

12. Integration of Credit Scoring and Portfolio Optimization

This work illustrates how wealth management is pioneered by fintechs that apply machine learning to improve the customer experience. An important offering in wealth management is the discretionary portfolio management in which the wealth manager, on behalf of the client, creates an investment portfolio and takes subsequent investment decisions and operation, without the client's involvement. An important element of this service is the compliance with the portfolio risk profile that is constructed based on the risk tolerance evaluation made by the expert in conjunction with the client. The principal mission of this work is to demonstrate how credit scoring and portfolio optimization decision-making can be made coherent with the client risk profile using modern machine learning and advanced numerical methods.

We start with the classical asset allocation problem of a risk-averse investor, exploring evidence corroborating the traditional optimization rule of quasi-concave utility function. We create a demonstration downloader that incorporates the utility function determined by a unique couple of risk aversion-event probability values. We then review multi-agent models of event occurrence probabilities predicted by the client sentiment indicator and machine learning classification models. We apply feedback Q-learning to enrich the set of probabilities triggered by the wealth manager. In the second part of the work we demonstrate how the user-specified risk profile is clustered and mapped into the credit scoring risk class values. These classes are obtained by the supervised forest tree regression model fitted on the training data created using loan portfolio performance annotated with multi-year client trade histories. Finally, we present packages and applications being developed for risk profiling and the dynamic evaluation of market event probabilities needed to comply portfolio disbursement and transaction limits specified by the wealth manager. The modularity of the packages allows their integration with other existing credit scoring and portfolio optimization solutions.

13. Regulatory Considerations

Financial regulations play a crucial role in protecting the interests of investors, depositors, and other participants in financial systems. They increase the demand for clarity and explainability of the financial institution policies behind model implementations. This comes in direct conflict with the demand for increasingly sophisticated machine learning models. To avoid being penalized for wrongful decisions, financial institutions are motivated to blindly follow rules such as maximizing model accuracy or minimizing model risk for deployment selection, which either result in models that provide zero value-add compared to classical models, or untested models that are costly to execute. Both alternatives go directly against the financial interests of such institutions.

A natural and exogenous for the development of machine learning models is the painstakingly detailed approach that must be taken to ensure that biased decisions are not made on protected or arbitrary variables such as race or ethnicity. In terms of implementation, the general philosophy is that monitoring and assessment must be performed to ensure performance longevity and forensics insights. This requires the performance of prescribed tests and procedures on the model before and after deployment. The result of this monitoring and assessment should subsequently be documented and reviewed regularly. These regulatory restrictions may directly impact the optimal selection of machine learning models and the interpretation of results. Therefore, regulators should be part of or at least consulted during development to allow for smoother pass rate assessments that consider the nuances of trained models.

13.1. Compliance with Financial Regulations

The notion of compliance with financial regulation serves to establish the boundaries of trust of the financial system that allow it to work smoothly and efficiently. A than undesirable situations to arise. These undesirable situations that can create significant losses and that led governments to intervene directly in markets around the globe, are: Predation, Lack of Efficiency, Distorted Market Signals, Negative Externalities, Economic Instability, Erosion of the Broad-Based Financial Contracts, Negative Externalities Cost, Financial Data Monopolization and Suboptimum Information Systems. In developed economies around the globe, financial regulations have been a consequence of past undesirable situations. The question here is whether or not utilizing machine learning and AI methodologies is, in itself, an undesirable situation? Moving further, is the utilization that at some stage being indicated by supervisory authorities as a-du-ponderable of the work performed by different firms, across desired main pillars of their exploitation or as transgressor of one or more prohibitive norms established in those pillars? Answering the first question, necessary should be a dialogue of proprieties with regulatory authorities to perhaps check if a group of different financial practitioners or monolithic executors in some markets can be considered transgressors or, on the contrary, motivated by employing a toolbox that delivers the best possible output for clients of the financial intermediators or accomplishment of financial interaction.

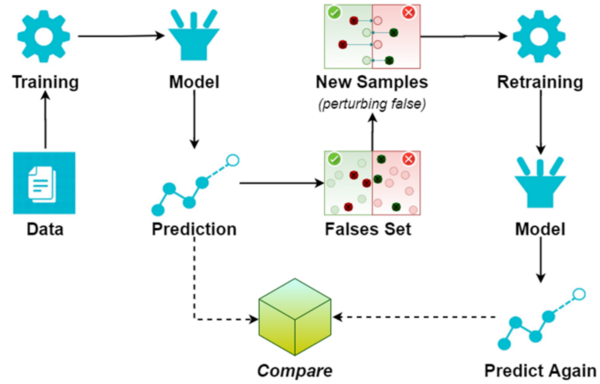

Fig : 5 augmentation strategy for credit scoring modeling

We have briefly touched on some regulatory norms that could somehow allow the approval of methodologies on a monitor and checking basis or that could illustrate the choice of adopting a portfolio of penalization or reward when addressing machine learning solutions. We have here regulatory principles that are behind the most significant basic areas of business. The overall goal of those norms is for all supervised firms to report results following accepted standards of truthfulness and discipline, which in turn would allow participants in the financial services ecosystem to trust the information received in their interactions with those firms.

13.2. Impact of Regulation on Machine Learning Models

Machine learning is a burgeoning area of research spanning various academic disciplines and scientific endeavors. As the focus of academic exploration continues to expand into broader areas, machine learning is becoming ubiquitous and is infiltrating many academic domains increasingly. However, it is not without its challenges and hurdles since deep-rooted beliefs and justifications supported traditional statistical techniques as being inherent modeling choices, while machine learning offers little explainability. For example, it is easy to compute the likelihood of data; provide a statistical interpretation for inference and prediction; and derive associated properties, but it is hard to justify for such explanations in machine learning-based models. As a result, statistical models in certain scientific domains have been able to maintain a stronghold over their machine learning counterparts.

Explanability is of paramount importance in our implementation for the various machine learning-based models, and justifications describing risk measures and credit risk are made succinctly. This is especially important since the financial aspect of credit risk has increasingly become the gauntlet against which machine learning, especially deep learning, receives its strongest pushback. For any implementation of models, we maintain that the regulatory considerations and implications need to be taken into account so that the hurdle in adopting such models is not insurmountable. In this chapter, we discuss the proposed emphasis by bank regulators for validating any proposed predictive model in finance for possible model risk. The kind of model risk and explanation required for structured finance models is far more stringent than for plain vanilla credit risk. Further, so-called deep learning models require specific methods for validation and risk quantification while using them for ongoing prediction during model draining.

14. Case Studies

In this chapter, we discuss a few successful implementations while providing additional supporting information to our work. We first review a wide selection of case studies describing some of the enterprise clients we encountered while providing consulting services to the wealth management industry. We explain to what extent and how much credit scoring and associated portfolio optimization efforts are presently being state of the art, the levels of automation and sophistication we have seen in the clients that we interviewed, and finally the industry standards. We then explain the lessons learned from these interviews. Here, our intent is to provide a view on what criteria clients and organizations have when selecting a firm to partner with for either or both of the components that constitute wealth management: portfolio construction and credit scoring.

The intention of our work was to contribute to providing wealth management companies a better understanding of how machine learning technology research could help them be more competitive in the future, and more specifically regarding their portfolio and credit scoring needs. To that extent, we set up a bunch of interviews with a wide selection of actors in this industry, from boutique firms to major players. We gathered insight on what technology these companies were currently using, what choices they made, what platforms they built their wealth management offers on, how satisfied they were with them, what were their clients' expectations about the quality and accuracy of the models that generated their financial investment recommendations, and how seriously disruptive they thought scientific consulting services were in these components' future. We ended up compiling a number of lessons learned to summarize our findings.

14.1. Successful Implementations

In the last several years, a number of new approaches to credit scoring and portfolio optimization have been implemented by fund management companies in Europe and the United States. The increasing competitiveness of the wealth management business has caused these companies to focus ever more attention to the issue of credit scoring accuracy. With the investment clientele broadening, including both investments and private banking of very affluent clients and investments of less affluent clients which tend to be socially less stable and reward less development for the bank, the business requires constant monitoring and periodic adaptation of the established credit scoring models. New neural network and support vector machines models have been introduced by a number of private banks for broad segments of affluent clientele. In addition, the derivative work has utilized decision trees and paradigm models based on other machine learning approaches have been successfully tested in this area. Experiments with addition of macroeconomic indicators have shown moderate improvement, but at the cost of model redundancy. Portfolio companies with manager skill observable in their long existing track records have been considered less risky by several banks implementing machine learning approaches, and models based on linear methods have also proven beneficial in combining the excess returns of such portfolio managers with their potential estimated realizations. Various dynamic techniques have been used for both asset allocation and dynamic redemption of funds. Research in the area of adaptive portfolio techniques utilizing supervised machine learning has pointed out to considerable potential of portfolio management in the field of hedge fund selection and investment.

14.2. Lessons Learned

Taking a step back and reflecting upon the use of machine learning techniques to solve real-world problems in wealth management, we highlight nine common pitfalls. These pitfalls provide a good description of useful lessons learned along the way and are essential to avoid, as they answer the "What could have gone better?" question. First, gathering high-quality, long, and granular datasets is a major problem in private wealth management. In most cases, public datasets don't exist, and acquiring or constructing high-quality internal datasets is a time-consuming and costly task. Having said that, if the organization already possesses credible datasets, it is usually worth to spend the required resources to analyze how the data can be used to improve existing methods and services. Second, data scientists should understand the business and why specific variables and data sources are important for wealth management. They should, ideally, work on the problem closely with wealth management practitioners. The involvement of wealth management professionals is key at multiple moments in the project, especially during the framing of the problem and the interpretation of results steps. At these moments, there is a risk of failing to properly account for wealth management characteristics, such as regulations, that have a major impact in the results. Machine learning scientists must not shy away from working closely with wealth management practitioners to minimize this risk. Third, whenever feasible, model simplicity should be valued as a model is easier to explain and interpret, has lower calibration costs, and is less sensitive to model risks. Fourth, the relatively small sample sizes present in wealth management could lead to the selection of overly complex models, overfitting, and worse results in production. Implemented models should then ideally be fine-tuned using real-world results from production. The performance of models built upon machine learning techniques should be monitored continuously. If possible, the models should be updated on a regular basis.

15. Future Trends in Machine Learning for Wealth Management

The opportunity of the digital economy, the profile of investors, the investors philosophy, the technical developments of AI contribute to the strong AI engagement in Wealth Management especially on customers relationship, investment banks, Fintechs and other stakeholders in the market apply AI on various use situations. With the continuously growth of the digital economy and the enhancement of decision-making unsupervised learning about the market in new emerging technology in AI, such as transformer and foundation model, the future of AI in Wealth Management could have a turn of it. The core value here is the “Plug and Play”. On many Wealth Management use situations with business scenarios and expertise knowledge heavily rely on, Widely speaking, the architecture leverage on the new technology could apply and work efficiently with professionals operators, and the closed local gap with automatic intelligent machines without or tiny supervisor works from knowledge workers. This would be realistic in Wealth Management for the high reward useful service with reasonable cost for asset management. The core idea of business models service these kind turn of AI architecture in the next stage of AI development is to give customers a better personalized experience but with transferable value in horizontal across customers in a specific use across domain specific market, and in particular is the super simple, super active, and semi-distance building of relationship between investors and investment units.

In the foreseeable future, AI will have more and more influence on financial markets as well as on working processes within the financial companies. Both will be supported by large language models, making the human-computer interaction easier, speeding up the process of data collection and data analysis, thus improving decision making. AI acts like an assistant without the limitations of skill and experience. The number of employees needed to supervise such processes will be much smaller than it is today. Nevertheless, such processes developed in cooperation with humans will increase efficiency and support decision-making in complex or less developed areas. In a longer perspective, AI will not only support the decision-making process, but by using its self-learning capabilities, it will also execute the processes behind it. In that sense, the employees will change from decision-makers to controllers.

15.1. Emerging Technologies

Rapid advances in technology, combined with shifting user behaviors and preferences, are driving pressures on financial services. As consumers in the digital age initiate business for all services via mobile, banks and wealth managers are striving to keep pace with the evolving digital landscape and investing heavily in digital marketing and customer acquisition. However, acquiring customers digitally, competing in a space offering attractive rates and fees, as well as the cost of serving them are putting significant pressure on profitability. Investing in digital distribution is essential for achieving scale and growth. However, wealth managers will also have to get much better at utilizing technology to help them provide superior service, advice and investment performance for client acquisition and retention. In the future, wealth managers are likely to serve both the high-net-worth and the mass affluent segments. This exclusive service, combining face-to-face high touch with online technology, presents real challenges. We suggest using machine learning models to support the advisors, augmenting their decision making and optimizing outcomes for the client. Such models can analyze a wide range of variables - the critical elements determining investment success - that would be beyond human ability. Over time, as algorithms are fed more and more data and learn from mistakes, the quality of their advice will improve. Machine learning will assist wealth managers in reducing costs, increasing productivity, and improving outcomes while creating scale. But using machine learning in optimization is only the start. However, currently only a fraction of advisors are making any use of machine learning or AI technologies to assist in investment decisions. No one wants to be the first to deploy in error. Nevertheless, machine learning and AI are here to stay. Our contention is that machine learning will make huge inroads.

15.2. The Role of AI in Finance

Artificial Intelligence (AI) has erupted on the scene, pushing technology and research to mostly unexplored avenues. AI is manifesting in various ways, through Machine and Deep Learning, Reinforcement Learning, Transfer Learning, Deep Reinforcement Learning, Text Mining, Neural Networks, Automated Planning, Game Playing, Expert Systems, Knowledge-Based Systems, Constraint Satisfaction, Reasoning Engines, Natural Language Processing, Computer Vision, Robotics, Search, and many more. What makes AI explode in acceptance and results is, one step after another, breakthroughs in new algorithms and implementations. Loading huge amounts of data from signal processing, dynamic programming, linear and non-linear optimizations, decision trees, Bayesian learning, or pattern recognition, not so new subjects, further enhanced by the capability of General Purpose Computing on Graphics Processing Units, by hardware and software renewals boosts power and efficiencies of such Deep Learning and Machine Learning techniques. AI applications excite everyone worldwide. Private life is cultivated through face recognition, specific music genres detection and suggestions, autonomous driving, natural language processing, social media findings and surfacing, internet searching, enhanced biometric security protocols, and many more. For corporations, tasks and processes in HR selection, recruitment and development, visualization of trends and customers' propensities, business and market developments and forecasts, behavior predicters and consolidators, and many others are being enabled by AI.

In the financial industry, results have arrived, are arriving, and will still arrive. Credit risk is being modeled based on Machine Learning and proposed alternative models outperformed the successful logistic regression. In-house research by huge investment banks used wealth of data, using Random Forest or lasso regularization provided warning signs in investors' portfolios before credit events, using classical but enhanced techniques. Asset management in large enterprises has been using optimization techniques, where asset returns could be modeled by Artificial Neural Networks, Decision Trees, Kernel Models, and tried to improve volatility through optimization.

16. Conclusion

The Banking sector regulates the credit system throughout the economy and allocates and disseminates the wealth, accelerating the pace of development for all provided that the rules are respected. With the introduction of Artificial Intelligence, new dimensions of technological leverage emerge at that time. The parallel involution of the pioneer Machines with the involution of the LAWS becomes an imperative in proposing a new "AI-based ordering of services". The AI gives the possibility to engage any user through a "market incentive" and to intercept the activity carried out by any user for any partner company. The introduction of Artificial Intelligence in financial activities, from Credit Scoring to Portfolio Management, allows to dramatically digitalize and derisk the Banking, Insurance, Finance sector, allowing to share knowledge for inclusion. It allows to make the transformation smart in that the Models made available at any time during the processing of the Request can be updated with the latest news, allowing the sensitive users to jointly solve the problem of the provision of consultancy service, to co-creating the shared value of Responsability, Trust and Loyalty with traditional dynamics. Finally, the deepening of Portfolio Analysis and the study of not only returns, but also of risks, damages and the consideration of Sustainability allow to go into to develop an ethical Finance where the value of the Capital is not only the useful but also above all the relationships and the consideration of Reference Stakeholders, possible and essential through the sharing of Knowledge.

8. References

1. Paleti, S., Singireddy, J., Dodda, A., Burugulla, J. K. R., & Challa, K. (2021). Innovative Financial Technologies: Strengthening Compliance, Secure Transactions, and Intelligent Advisory Systems Through AI-Driven Automation and Scalable Data Architectures. Secure Transactions, and Intelligent Advisory Systems Through AI-Driven Automation and Scalable Data Architectures (December 27, 2021).
2. Gadi, A. L., Kannan, S., Nanan, B. P., Komaragiri, V. B., & Singireddy, S. (2021). Advanced Computational Technologies in Vehicle Production, Digital Connectivity, and Sustainable Transportation: Innovations in Intelligent Systems, Eco-Friendly Manufacturing, and Financial Optimization. Universal Journal of Finance and Economics, 1(1), 87-100.
3. Someshwar Mashetty. (2020). Affordable Housing Through Smart Mortgage Financing: Technology, Analytics, And Innovation. International Journal on Recent and Innovation Trends in Computing and Communication, 8(12), 99–110. Retrieved from https://ijritcc.org/index.php/ijritcc/article/view/11581.
4. Sriram, H. K., ADUSUPALLI, B., & Malempati, M. (2021). Revolutionizing Risk Assessment and Financial Ecosystems with Smart Automation, Secure Digital Solutions, and Advanced Analytical Frameworks.
5. Chava, K., Chakilam, C., Suura, S. R., & Recharla, M. (2021). Advancing Healthcare Innovation in 2021: Integrating AI, Digital Health Technologies, and Precision Medicine for Improved Patient Outcomes. Global Journal of Medical Case Reports, 1(1), 29-41.
6. Just-in-Time Inventory Management Using Reinforcement Learning in Automotive Supply Chains. (2021). International Journal of Engineering and Computer Science, 10(12), 25586-25605. https://doi.org/10.18535/ijecs.v10i12.4666
7. Koppolu, H. K. R. (2021). Leveraging 5G Services for Next-Generation Telecom and Media Innovation. International Journal of Scientific Research and Modern Technology, 89–106. https://doi.org/10.38124/ijsrmt.v1i12.472
8. Adusupalli, B., Singireddy, S., Sriram, H. K., Kaulwar, P. K., & Malempati, M. (2021). Revolutionizing Risk Assessment and Financial Ecosystems with Smart Automation, Secure Digital Solutions, and Advanced Analytical Frameworks. Universal Journal of Finance and Economics, 1(1), 101-122.
9. Karthik Chava, "Machine Learning in Modern Healthcare: Leveraging Big Data

for Early Disease Detection and Patient Monitoring", International Journal of Science and Research (IJSR), Volume 9 Issue 12, December 2020, pp. 1899-1910, https://www.ijsr.net/getabstract.php?paperid=SR201212164722, DOI: https://www.doi.org/10.21275/SR201212164722
10. AI-Based Financial Advisory Systems: Revolutionizing Personalized Investment Strategies. (2021). International Journal of Engineering and Computer Science, 10(12). https://doi.org/10.18535/ijecs.v10i12.4655
11. Cloud Native Architecture for Scalable Fintech Applications with Real Time Payments. (2021). International Journal of Engineering and Computer Science, 10(12), 25501-25515. https://doi.org/10.18535/ijecs.v10i12.4654
12. Innovations in Spinal Muscular Atrophy: From Gene Therapy to Disease-Modifying Treatments. (2021). International Journal of Engineering and Computer Science, 10(12), 25531-25551. https://doi.org/10.18535/ijecs.v10i12.4659
13. Pallav Kumar Kaulwar. (2021). From Code to Counsel: Deep Learning and Data Engineering Synergy for Intelligent Tax Strategy Generation. Journal of International Crisis and Risk Communication Research , 1–20. Retrieved from https://jicrcr.com/index.php/jicrcr/article/view/2967
14. Raviteja Meda. (2021). Machine Learning-Based Color Recommendation Engines for Enhanced Customer Personalization. Journal of International Crisis and Risk Communication Research , 124–140. Retrieved from https://jicrcr.com/index.php/jicrcr/article/view/3018
15. Nuka, S. T., Annapareddy, V. N., Koppolu, H. K. R., & Kannan, S. (2021). Advancements in Smart Medical and Industrial Devices: Enhancing Efficiency and Connectivity with High-Speed Telecom Networks. Open Journal of Medical Sciences, 1(1), 55-72.
16. Chava, K., Chakilam, C., Suura, S. R., & Recharla, M. (2021). Advancing Healthcare Innovation in 2021: Integrating AI, Digital Health Technologies, and Precision Medicine for Improved Patient Outcomes. Global Journal of Medical Case Reports, 1(1), 29-41.
17. Kannan, S., Gadi, A. L., Preethish Nanan, B., & Kommaragiri, V. B. (2021). Advanced Computational Technologies in Vehicle Production, Digital Connectivity, and Sustainable Transportation: Innovations in Intelligent Systems, Eco-Friendly Manufacturing, and Financial Optimization.
18. Implementing Infrastructure-as-Code for Telecom Networks: Challenges and Best Practices for Scalable Service Orchestration. (2021). International Journal of Engineering and Computer Science, 10(12), 25631-25650. https://doi.org/10.18535/ijecs.v10i12.4671
19. Srinivasa Rao Challa. (2021). From Data to Decisions: Leveraging Machine Learning and Cloud Computing in Modern Wealth Management. Journal of International Crisis and Risk Communication Research , 102–123. Retrieved from https://jicrcr.com/index.php/jicrcr/article/view/3017
20. Paleti, S. (2021). Cognitive Core Banking: A Data-Engineered, AI-Infused Architecture for Proactive Risk Compliance Management. AI-Infused Architecture for Proactive Risk Compliance Management (December 21, 2021).
21. Vamsee Pamisetty. (2020). Optimizing Tax Compliance and Fraud Prevention through Intelligent Systems: The Role of Technology in Public Finance Innovation.

International Journal on Recent and Innovation Trends in Computing and Communication, 8(12), 111–127. Retrieved from https://ijritcc.org/index.php/ijritcc/article/view/11582

22. Venkata Bhardwaj Komaragiri. (2021). Machine Learning Models for Predictive Maintenance and Performance Optimization in Telecom Infrastructure. Journal of International Crisis and Risk Communication Research , 141–167. Retrieved from https://jicrcr.com/index.php/jicrcr/article/view/3019
23. Transforming Renewable Energy and Educational Technologies Through AI, Machine Learning, Big Data Analytics, and Cloud-Based IT Integrations. (2021). International Journal of Engineering and Computer Science, 10(12), 25572-25585. https://doi.org/10.18535/ijecs.v10i12.4665
24. Kommaragiri, V. B. (2021). Enhancing Telecom Security Through Big Data Analytics and Cloud-Based Threat Intelligence. Available at SSRN 5240140.
25. Rao Suura, S. (2021). Personalized Health Care Decisions Powered By Big Data And Generative Artificial Intelligence In Genomic Diagnostics. Journal of Survey in Fisheries Sciences. https://doi.org/10.53555/sfs.v7i3.3558
26. Data Engineering Architectures for Real-Time Quality Monitoring in Paint Production Lines. (2020). International Journal of Engineering and Computer Science, 9(12), 25289-25303. https://doi.org/10.18535/ijecs.v9i12.4587