Deep Science Publishing
https://doi.org/10.70593/978-93-49910-98-0_4

Chapter 4: Machine learning for predictive analysis in market forecasting, credit scoring, and customer behavior

4.1. Introduction

Predictive analytics refers to the use of predictive models, cutting-edge statistical techniques, and machine learning applications to estimate by giving predictions for future data and events. Such techniques are important components of analytics and report explicitly in overcoming the limits of traditional uses of data analysis in the planning and execution of markets and business functions in providing desired performance. Predictive models are based on estimation, calibration, and validation by applying mathematical and statistical relationships between available data (Khan & Singh, 2023; Das, 2024; Jain & Sharma, 2024). These models are used to predict prospect information such as customer preferences, attitudes, interests, and buying specifications or needs; market trends and other characteristics of market supply and demand; and business performance in such areas as sales revenue, profitability, market share, customer retention, and customer lifetime value. Predictive analytics seeks ways to project such behaviors and outcomes into the future. It provides structure and systemization that methods of predictive modeling apply to use for other and often much less difficult-to-quantify, practice, and art forms of computing the correct answer. Over the years, businesses have sought difficult-to-acquire and often expensive enterprise data like customer contact information. By investing in data infrastructure development, data analytics, and data mining and employing capable data information experts, with data and business knowledge, and practitioners, data have been successfully converted into business, employee, and customer values. Machine learning algorithms and applications, together with appropriate data, have been applied to complete complex tasks requiring optimal or near-optimal performance. Well-trained and appropriately applied, these methods have achieved higher performance objectives and at lower costs than traditional methods (Li et al., 2024; Zhang et al., 2025).

4.2. Overview of Machine Learning

Machine learning is a branch of computer science that uses algorithms to parse data, learn from it, and then make a decision or prediction about something in the world. The algorithms work by analyzing training data to build a mathematical model, then applying that model to new data to make a prediction. In other words, machine learning allows computers to learn on their own without being specifically programmed for each task. Machine learning tasks are usually classified into three categories: supervised learning, unsupervised learning, and reinforcement learning. Supervised learning is where we provide a set of input-output pairs, and the model learns a mapping from inputs to outputs. The two main types of supervised learning problems are regression, where the output value is a real number; and classification, where the output value is a member of a discrete set. For example, we could use supervised learning to predict the price of a house, or predict whether the email is spam or not. The second task, called unsupervised learning, is where we provide only input data and the model discovers any patterns present in the data. A common unsupervised learning task is clustering, where we want the model to group the input data into clusters that are similar to each other. Reinforcement learning is a different kind of machine learning task. It is inspired by behavioral psychology. Here, an agent learns to interact with the environment in order to maximize some notion of cumulative reward. Unlike the other two types of machine learning, the agent is not used to make predictions about the environment. Rather, it learns to take actions in the environment to achieve some desired goal. Reinforcement learning is commonly applied to robot control, game playing, and any domain where an agent makes a sequence of decisions.

4.3. Predictive Analysis in Market Forecasting

In this more technical domain of supervised learning the goal is to predict the future behavior of an uncertain time series from current or past information about it. The dynamical models that are normally used to make these predictions are called forecasts. Forecasting has a long and rich history, having first been performed by ancient astronomers, and employed since the early beginnings of modern statistics. It is possible to forecast numerous diverse phenomena ranging from the size of street crowds to stock prices, and the number of births in a month. Dealing with uncertain data is admittedly difficult, often around 90% of the efforts put into data mining projects are spent pre-processing the data to make it ready for the prediction process.

Most prediction methods offered in prediction software are specific to a certain domain and designed by experts in that domain. They include methods such as ARIMA modeling which is widely used in finance, engineering and economics for the prediction of univariate time series data. In recent years, scientists from diverse backgrounds have begun to collaborate and invent more universal prediction methods that have been shown to work well in many different domains. These methods are much easier to use, and often more accurate too. One of these new model families is the ensemble methods. Ensemble methods have been highly successful at winning prediction competitions in many domains. They have obtained sweeping victories on the monthly forecasting competitions, and have also been proven to work well for the prediction of univariate time series data in the next months of the months in the Global Temperature database.

Fig 1 : Predictive Analysis in Market Forecasting

4.3.1. Importance of Market Forecasting

Market forecasting uses historical and contextual data to make future projections. In data-driven business analytics, marketers use forecasting to predict future sales across time and geography to assess plan efficacy and product competitiveness and to set budgeting and supply chain needs. Forecasting models can take many forms, using techniques such as time series analysis or chain ratio method. Time series analysis uses current and prior period actual figures to project future sales. This is a common technique given the relative ease of access to data. Time series analysis is useful for detecting seasonality in data sets and can take the form of a seasonal or working day method for projections. Chain ratio methods use economic, demographic and market variables to determine future sales projections and require a more complex model of historical customer behaviors to calculate the ratios applied to future projections. Given the stakes, be it company budgets or investment decisions, accurate forecasting is vital. Predictive analytics can create objective rulesets based on explicit constraints on sales or customer behavior that can be used to make these forecasts. Nowadays, forecasting is an afterthought and considered a simple process, when it should be our most important task to improve profitability.

Forecasting can be one of the most important and valuable tasks provided by machine learning. Accurately predicting future behavior for almost any area of business would be incredibly valuable. With better accuracy in forecasting, budgets can be better allocated to the resources to ensure better SLA adherence. This has an obvious revenue impact on the business. Budgets can be cut to enable the business to be competitively positioned for any downturns. Customers can be sent targeted offers to maximize the likelihood of impact. These predictions are not necessarily service demand forecasts but could cover any area of business interest.

4.3.2. Machine Learning Techniques in Forecasting

Forecasting problems involve modeling underlying patterns from time-series data in order to predict their future values. The focus of the forecasting process is on predicting what happens in the future rather than on elucidating how and why the proposed model works. Analyzing historical data can only partially provide the knowledge needed for making good forecasts. Prediction per se involves a set of methodological tools that has been developed independently of the basic mechanisms of the system being forecasted. The process of forecasting involves using a model of the system being studied to generate a value for the dependent variable(s) at the forecasting point(s). This prediction is then adjusted for its expected bias, anticipating the fact that such a model will not generate predictions of zero bias.

The forecast error is defined as the difference between the actual value predicted and the value obtained with a prediction model. It is seldom that the prediction error models are independent of the forecasting process either in the analytical or in the empirical formulation. In fact, these error-variable relationship models are usually the basis for calculating the measure of efficiency. In most realistic forecasting situations, a prediction is based on the knowledge available from a multivariate history. Without special conditions, each variable has its own prediction. In addition to classical multivariate forecasting methods, machine learning may incorporate explicit physical relationships. While many modeling techniques can be applied to the probability density function of the process involved, traditional inference is based on an empirical joint distribution. It therefore embeds the dependency structure. Several techniques produce probabilities under certain conditions, but having the conditional probabilities makes it especially feasible to explore other possible scenarios.

4.3.3. Case Studies in Market Forecasting

Kourentzes and Wu employed second-level machine learning ensembles models in real-world demand forecast scenarios. They concluded that their proposed model is statistically better than the benchmark and use it to demonstrate the potential economic impact on a company's profits. Kourentzes and Makridakis focused on automatic selection of exogenous variables and they showed better performance than the economic experts. The authors employed univariate time series, with ARIMA and exponential smoothing methods as reference models. Kang and Yeom used a univariate time series approach to optimize the model choice, data choice, and parameter choice in sales forecasts of the Korean companies from various sectors. The authors focused more on fault tolerance than prediction accuracy. Their results showed that, based on the median MaSE performance, the multilayer perceptron and the model learned using the bobyqa algorithm were the best overall model.

Fildes and Makridakis argue that "the methods applied should take advantage of previous experience and understanding of the way stochastic processes governing the demand work" and "the available forecasting methods should then be used in an optimal way." The forecasting process cannot be reduced to merely applying a machine learning algorithm to forecast, and demand forecasting is a "backed-in-process." The authors compare different forecast procedures that embed machine learning models by means of accuracy measures and proposed a two-phase approach. They demonstrated the importance of a sound data preparation phase to improve forecasting accuracy as well as the methods' performances. The first three months of the quarterly data are used for comparison and forecasting phase; the authors employed machine learning methods implemented in different packages. Their results showed original simple exponential smoothing generated the best forecasts, although many other methods are almost as efficient.

4.4. Credit Scoring with Machine Learning

Credit scoring models predict the likelihood of a loan being defaulted. There are huge economic and social implications stemming from these predictions of default probability: banks rely heavily on credit scoring models, and generally need to be conservative with their estimations. However, it costs money for consumers who would not yet default to pay higher interest rates due to mispredicted default probabilities. These predicted probabilities are also important in many risk-analysis related applications such as the pricing of government backed loans – here, mispredictions can lead to government bailouts in case of general economic downturn. The opposite is the example of institutions that do not offer safe loans, and are facing the chance of declaring bankruptcy in case of a general downturn. Again, underestimating default probabilities might lead to economic crises where public funds are needed to cover the institutional default. Given the probabilistic nature of the predictions, the evaluation of credit scoring models needs to be carefully done, and is often performed with a host of different criteria. For this reason, this task in particular is well-suited for machine learning: the error on many of these metrics can be minimized making combination predictions and stack ensembling easy. However, the problem is still difficult because of the extremely imbalanced nature of default to non-default cases.

The traditional methods use logistic regressions, amongst others survival analysis, decision trees, ensemble methods such as random forests or boosted trees, and recently neural networks. Banks appreciate the interpretability of the model. The traditional models are often not performing terribly well either: One study comparing the logistic regression model to random forests found that increasing the area under the curve potential from a mean improvement of 8% provides an economic benefit of up to USD 35 billion. Time horizon studies have shown that conventional logistic models tend to perform worse over a longer period of time.

4.4.1. Traditional Credit Scoring Models

Machine learning generates added value in investment decisions. Credit scoring applications, particularly in the retail sector, have had more than 50 years of success using credit scoring models to guide and govern decisions on credit acceptance and loans. Automating these predictions allowed new financial institutions to enter the market with low levels of initial investment. From a client's perspective, these vending machines created a frictionless process, allowing customers to receive real-time decisions on credit requests without any human involvement and at very low costs. Credit scoring models favor well-collateralized credit operations. Regardless of their location, retail customers are all considered individual origins of risk in a homogeneous pool. These models are based on discriminatory analysis that classifies customers into two classes: those who will reimburse the funding and those who will default. The main hypothesis is that their articulated characteristics are of a similar nature. Customers are identified using "hard" data describing their background or track record, including age, marital status, number of children, address, and way of life, and "soft" data regarding their credit behavior, which mainly includes previous operations. Logit analysis, probit analysis, and decision trees are the most popular methods. They seek to find the best linear combination of variables that offers the best fit for the data sample. This would be related to the probability of reimbursement. Credit decisions are made according to a threshold level, below which definitively reject the request. The important issue is the decision you would make with a new client that doesn't belong to the database used for calibration. The rules or proportions of the model are constant. For approximately 50 years, more than 80% of the decisions have been made according to heuristic rules of thumb. These models have already been criticized in the marketing decisions associated with credit shopping.

4.4.2. Machine Learning Approaches

So, the history of credit scoring, in a nutshell, is that financial institutions traditionally relied on statistical methods to predict creditworthiness. The advent of machine learning introduced several advanced techniques that outperform traditional models for this task. In recent times, there has been growing interest in the application of machine learning to credit scoring, as evidenced by a plethora of proposed techniques. Consider how most of the machine learning approaches were proposed relatively recently, especially when compared to credit scoring that dates back to the early 1930s, not to mention statistic models that preceded it. That is not to say that the entire area is saturated. In fact, it is generally accepted that it would be beneficial to further explore specific adaptations of some popular machine learning techniques, variants that could exploit available domain knowledge or semi-supervised and unsupervised learning techniques that could exploit unlabeled examples, ensembles, or combinations of deep learning with boosting and semi-supervised learning algorithms. The main issue is whether there is enough room for new proposed adaptations to warrant the issues associated with credit scoring, such as the ones relating to the regulatory requirements or data privacy. Unfortunately, one of the implications of the lack of interpretability associated with most machine learning techniques is that organizations must consider even more these issues. In summary, it would appear that the main road forward proposed to pioneer researchers would be to explore present machine approaches, in particular those with a proven track.

4.4.3. Impact on Financial Institutions

There is a shift in the banking sector towards credit risk prediction using advanced machine learning algorithms. These algorithms improve decision-making processes and resulting gains and lower risks. Consequently, owing to the fast growth in big data applications in financial systems, the focus of banks is moving from classical statistical-based models to machine learning data-driven ones. Banks can tell which customer is more likely to default based on long lists of historical, demographic, and other data with higher predictive power by using the more powerful machine learning models. These models can sift through multiple layers of complex and non-linear interactive relations behind the probability of defaulting, and help infer insights that top management can act on for better credit risk prediction and other business decisions.

The machine learning methods used by banks can help infer sensitive behaviors of investors, and present them with targeted solutions. Financial institutions can take proactive steps to identify customers who are likely to be dissatisfied with the service and who have alternative options. With credit rating predictions, the banks can either cut down on lending to an emerging market or increase interest rates of borrowers who are less likely to default. This truly poses problems in the industry as banks that engage in dynamic and sophisticated prediction strategies would face lower risks than those that do not. As these diversionary consequences of predictive behavior come into play, good decisions taken on predictions make these advanced methods of prediction even more pointed. Through inferential strategies, by making efforts in allocating resources in a more effective way, banks can earn higher profits relative to costs as compared to those banks that take the classical decisions that are not driven by predictions.

4.5. Understanding Customer Behavior

Customer behavior plays a critical role in determining a company's success. Business activities that are aimed at understanding customer behavior are the foundation for effective corporate strategy including new product development and positioning, market communication, distribution strategy, and channel relationships. Analyzing customer preferences and behavior helps marketers recognize customer groups that are interested in a specific product. In particular, market basket analysis, which builds models based on groups of products that tend to be purchased together, helps determine prices for co-depend products. Customer micro-segmentation allows a business to understand its own internal-product relationships as well as determine preferred supplier profiles. This is especially useful in B2B markets where there are few suppliers. Will the customer preference for cheaper products and switching to another supplier in the class attribute will be high? If we know, we can target specific customers in an aggregated way rather than on an individual basis.

Customer segmentation is any division of customers into groups; generally, we want homogeneous groups with respect to some quantitative measure, such as sales or satisfaction index. We want to study customers – either individually or as segments – based on their profitability. Based on distances, clustering algorithms create and grow partitions that define clusters. The most prominent clustering algorithms use Euclidean or other distances among observations, but other approaches, such as market basket analysis, are based on conditional probabilities of products being bought together. Association rule algorithms detect transactions in a database that contain reflecting items with a support index that is above a certain level. Clustering customers via RFM or holding periods can help identify those who contribute to a bank's profitability, although behavioral segmentation based on clusters of customers can be more powerful. Using propensity models, customer segments can be further defined by cutoff models, often invoked for special offers.

4.5.1. Customer Segmentation Techniques

Customer segmentation is essential in the domain of understanding customer behavior. Oftentimes, companies have diverse customer bases. As a result, these organizations offer different customer mix and merchandise services. Importantly, in order to avoid cannibalizing higher margin customer, market segmentation should provide a concentrative research and analysis on key group of customers. Furthermore, customers' needs and wants are diverse creating different business values for organizations. Identification of groups of customers thus helps firms focus scarce resources in creating more effective marketing programs. The marketing idea of segmentation has existed for many decades.

Market segmentation can be defined from a marketing research perspective as the division of a market into relatively homogeneous segments having similar needs for rights and services. Businesses can then develop and communicate meaningful differentiative messages to each valuable market segment. In order to do this successfully, segmenting characters are critical. Traditionally usage based variables such as customer brand loyalty and frequency of use are used. However, it is difficult to establish the connection between these segmentation variables and actual purchase behavior. In more recent years, variables such as promotional response propensity, customer lifetime value, long-time switcher status, and the number of value added features are used. These components measure the post transaction behavior motivations and performance of consumers which are more closely related to actual business results.

Using machine learning, the traditional methods of segmentation such as demographics, firmographics, geo-demographics, lifestyle, and behavior can be significantly improved, enhanced, and complemented with advanced multivariate analysis. These algorithms in fact are able to analyze complex interaction effects and generate huge amounts of variables in person centric statistical models. These usage based predictor variables increase the predictive power, cost-effectiveness, and efficiency of businesses segmentation. Such diverse segmentation methods provide better customer profiles resulting in more timely and effective marketing decisions.

4.5.2. Predictive Modeling for Customer Behavior

In addition to inferring customers' attributes or affinities, predictive modeling techniques can be used to predict customer behavior explicitly in several contexts. The types of behavioral predictions that advertising campaign managers are the most interested in modeling are customer response to marketing communication efforts, financial transaction defaults, customers' people-based behaviors, and privacy. In the marketing context of customer response modeling, marketers are interested in modeling which customers will respond to a marketing offer, such as a campaign for a sale or a new service. Since the goals of these kinds of marketing operations are the generation of sales, ongoing relationships, and goodwill from customers, the modeling of customer

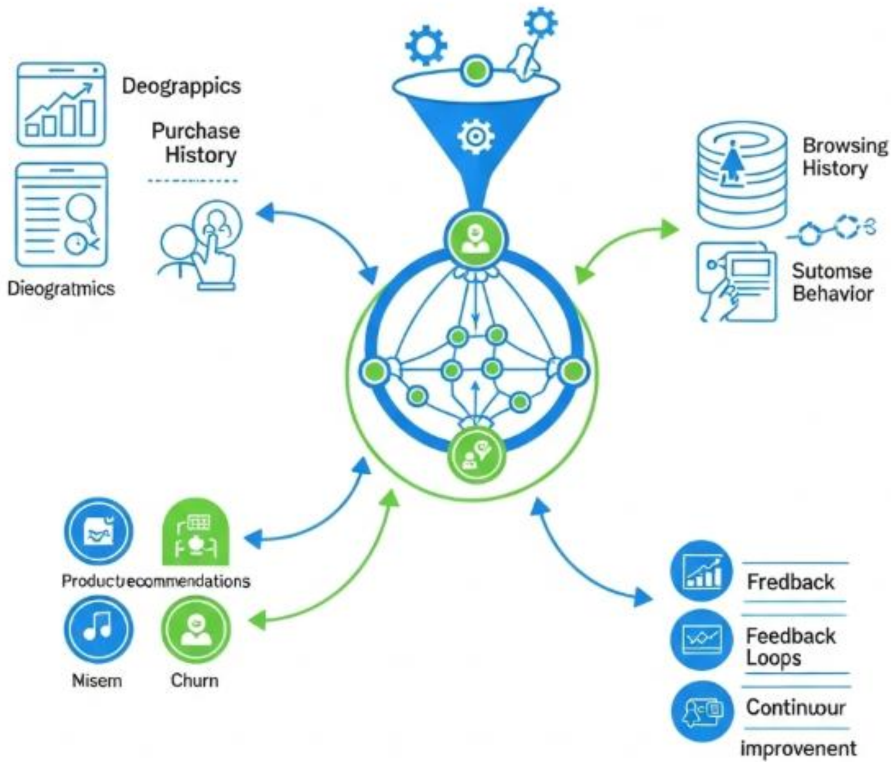

Fig 2 : Predictive Modeling for Customer Behavior

response has, and is likely to continue to have, a critical role to play in marketing.

Several marketing activities can benefit from customer response modeling efforts. For example, response modeling can be used to help firms select which customers should be included in a marketing campaign designed to generate immediate purchases. In this case, a firm tries to identify a group of customers who are likely to respond positively to the campaign offer. The firm then targets this group during the campaign. The firm's initial goal is to increase the sales generated by the promotional campaign to maximize return on investment. By identifying the customers likely to respond positively to the promotion, the firm can better control the expected campaign costs and thus improve the effectiveness and efficiency of the activity.

With the advances in data and computational techniques, customer response modeling is no longer limited to binary responses but can predict the full distribution of responses that customers are likely to generate during a campaign. Specifically, firms want to know three aspects of the expected customer responses: the number of response transactions, the time of these transactions, and their value. For the majority of customer response models currently used in marketing practice, estimates are obtained mostly for the expected number of transactions within the time horizon of interest, with the average transaction amount conditional on being positive being taken from a separate model. In most cases, the response models estimate the expected number of response purchases during the campaign time period. Hence, the response forecasts from the models can be viewed as synthetic customer purchase data for the campaign, with an expected number of purchases per customer during the campaign that is distributional rather than point forecast.

4.5.3. Applications in Marketing Strategies

The effectiveness of marketing strategies is evaluated basically through incrementing market shares, sales and maximizing profits at appointed levels of expenditures. In the past few years, traditional techniques of these equitable marketing strategies have been improved with the efficiency of a manageable benchmark. Effectively, a good benchmark or comparison which is about performance among old or in case new products that are similar, would give the distinguishable advantage of improving performance of the product or seeking what is preventing higher performance in allocating segments. The introduction and potent increase in the usage of Neural Networks have somewhat transformed the Machine Learning dimensions, such as information, quantification of certain equations and computing capacities.

In a range of from ubiquitous marketing to heavy marketing there may be producers or makers of these brands or products who have a potential to understand such networks. The keys to success on position and print assortment are "everyone is different," "audience complexity is growing," and "a good deal is no longer good enough." These optimistic slogans urge the marketers or strategists to look at the specific use of Marketing or Advertising Media to achieve their stated or considered aim or outcome. Along with other tools of the information, distribution and sales management, marketing or advertising and its slightly expanded version marketing communication remain the little core of most companies marketing strategy. It is said that these optimist slogans urge the marketing strategists and the groups to consider Product Attributes or Characteristics. Service Assembly Capacity and Cost, Scale of Resources or Inputs, Marketing or Product Communication Programs, Consumer Preferences and Behavior Patterns, Competition Marketing Strategies, and Available Production Capacity related to the predicted demands to marketing strategy decisions making on two and three times full-order potential.

4.6. Data Collection and Preparation

The foundation of producing useful ML models is data collection and preparation. It is a critical yet often overlooked step that can introduce bias and artifacts into an ML model. The first step in developing an ML model is to collect data that is related to the outcome. For example, if the goal is building a model to forecast company profits, the related data include historical company profits, macoeconomic variables, stock price data, and so on. It is best to use multiple sources of data because different types of data can introduce synergy effects to improve the ML model's forecasting performance. For example, the ML model can be made more accurate with news article data or stock price return data than with numerical time series data alone. This is true as long as multiple sources of data maintain similar temporal proportions. However, data from different sources may use different sampling frequencies. For example, financial reports are typically released quarterly by a company, while macroeconomic variables are released monthly by the government. In this case, the researcher can have the ML model train with the available data and see if that is sufficient. In most forecasting problems, temporal data is usually used as the predictor variable, while the label of interest is the data from another source, such as stock prices or companies' quarterly profitability.

In general, there are two types of data, structured and unstructured. Structured data includes numerical time series data and categorical data. In contrast, unstructured data can be textual data from news releases or qualitative data that needs to be converted into categorical or numerical format. For numerical data, cleaning outliers can smooth unexpected blip-like changes in the data and improve the ML model's performance. Additionally, all numerical data must share the same sampling interval or frequency along the time axis. Categorical data from financial databases typically require data transformations to represent the labels in numeric format. For textual data, either a cleaned or tokenized method should be used. The choice of which method to use is highly situation-dependent, as both methods have shown to work well.

4.6.1. Sources of Data

The data used for forecasting and customer insights generally comes in two forms. The first is the classical economic and social factors introducing into econometric or statistical models, which traditionally had to be collected manually from a variety of sources on a yearly or more often quarterly basis. The second is the new data, which in recent years has proliferated and from which a number of additional insights can be generated for forecasting purposes. New data, unlike the classical data, is not necessarily collected in a standard format. It can be short text information extracted from communication, product reviews, click-through on the stock or service, or real estate transactions scrapped from various sources. A particular focus of this volume lies in the use of alternative data and how scholars and practitioners have exploited this new data.

Classic economic data are more readily available for economies and societies at an aggregate level, while the alternative, new data can usually provide insights at the disaggregate, regional, functional, or even firm level, as they originate in specific segments or parts of the economy that may show signs of trends before the economy as a whole starts to show similar indications. As such, alternative nonstandard data could be seen to provide economists and business practitioners with a new tool to examine questions that could not be targeted before, or in areas for which reliable data were seldom available. In particular for market forecasting, the growth of high-frequency variation in financial as well as product market variables has led researchers and practitioners alike to use data with a similar frequency for their analysis.

4.6.2. Data Cleaning and Preprocessing

After data extraction, we systematically curate and prepare the data before analysis. This process requires careful consideration and skill because the performance of our methods, ultimately, is highly dependent on the quality of the processed data. Data cleaning refers to the process of preparing the data by correcting inaccurate, incomplete, or otherwise problematic data points. For instance, we deal with empty cells found in the data with below-average stock price or transaction volume for a particular firm or day. These empty cells can occur very frequently in financial datasets that need to be filled. Including these empty cells can bias the results, so we opt to test different methods of filling them. We also need to make sure that the prices of stocks are not negative during the transaction period.

Data preprocessing means transforming cleaned data into a suitable format that can be ingested by machine learning algorithms. Typical preprocessing steps include feature alignment, selection, and extraction. In the context of time series analysis, feature alignment refers to the synchronized mapping of one or more input variables with the same-sized output. During the period of our market reaction analysis, there are some days in which we do not have the same number of unique terms for our sentiment proxies. Consequently, we opt to merge our input features column-wise, based on day. Such a merge operation results in a dataset where each row consists of unique terms for a day, and when we train our models, they would be fed with a vocabulary that is representative of the target date's market sentiment. Since daily unique term counts can be quite diverse, we create a variable-length vocabulary for each day. In order to predict future stock price movement, we ensure the model input consists of previous days' sentiments.

4.7. Feature Engineering

An important aspect of any data science project is feature engineering, which is the process of selecting the most relevant properties of the data to use as input to the machine learning model. Characteristics can be chosen from the existing data or created from unavailable data using other methods, including techniques from signal processing, statistics, or domain knowledge. This is an important aspect of every machine learning project because the quality of the input data used is key for the predictive power of the model. More data is not always better, especially if more data correspond to irrelevant features. The task of feature engineering is therefore to identify and create relevant and informative properties that will make accurate predictions possible. Feature engineering for market forecasting aims to transform raw data into signals that make forecasting accurate and efficient. The design and selection of features should be strongly dictated by the underlying structural assumptions regarding the data-generating process.

Importance of Feature Selection

The relevance of feature engineering has been widely recognized. The single most important thought in machine learning research is that different learning algorithms may behave very differently on the same problem, with error rates that can differ by orders of magnitude. In many cases this is due to bad feature choice, and feature engineering remains a black art, relying heavily on human intuition, experience, and ingenuity. Feature engineering is crucial to success. If you want to build a model with hyper-parameters that are not too sensitive and which generalizes well to unseen examples you need to add relevant characteristics. If many irrelevant characteristics are available you should remove them or make sure that your method can handle it. Inadequate feature choice leads to models with over-fitting problems – models sensitive to perturbations of the training data – or under-fitting problems, where the model cannot explain any of the training data. Different types of models require different checks, and will be more or less sensitive to problems coming from feature overloading or restrictions.

4.7.1. Importance of Feature Selection

The essential machine learning requirement for building successful predictive models is to provide informative features that help techniques produce very accurate predictive models. It has been discovered that more than 80% of the machine learning modeling efforts are devoted to feature engineering. Therefore the right features not only enable a variety of learning techniques to achieve very high accuracy, but they also reduce the complexity of the models, improve the speed of execution for inference and training, and reduce the storage requirements for the model.

In market forecasting applications, the most widely used historical data augmented by a limited number of additional information such as weather, holiday seasons, and various forms of economic indices have been used for feature generation. For deriving demand for new products, usually the work is focused on deriving domain-specific features that could enhance the predictive accuracy of demand for the specific product. These domain-specific features in marketing are derived from the knowledge of the new product in comparison with similar or related products that have been launched in the past and the available demographic information for the units affected by the launch of the new product. The domain-specific feature creation is a process of exploring the previously available data for the similarities and differences for the products which have already been launched and using these insights to augment the new product launch features. The effectiveness of these domain-specific features depends on the accuracy of discovering new product related products from the available historical products' launch or impact data. Such domain-specific features help improve predictive accuracy only for specific product launches and require expert input for choosing the suitable features.

4.7.2. Techniques for Feature Engineering

The two basic tasks in feature construction are feature selection and transformation. Feature selection is to find a suitable subset of existing features from the original feature set. The selected features should help improve the performance of the supervised learning algorithm and eliminate noise for unsupervised learning. Feature transformation, on the other hand, should create additional features from existing ones. Both tasks aim to create a new feature subset that improves learning performance. Feature selection can be filtered or wrapped into the two types. Filter methods usually choose a subset of features according to heuristics that depend on the input data structure. Many heuristics have been proposed, and four common classes are: distance measures, dependency measures, information theory measures, and others. Wrapper methods usually optimize the candidate feature set by testing the supervised learning algorithm on the validation data set. Since a supervised learning algorithm is very expensive to test, the optimization is usually based on hill climbing or genetic algorithms. Wrapper methods are usually better than filter methods in that they can search for a suitable feature subset with respect to the supervised learning algorithm. However, they are much more computationally expensive than filter methods. Hybrid methods combine both techniques to leverage their advantages. Feature transformation methods transform the original features to a new space. In the transformed space, the features become orthogonal to each other. Some common techniques include kernel methods, basis expansion, change of variables, data projections, and local representations. Kernel methods require the algorithms to compute the similarity between any two data points without explicitly mapping the input data into the transformed space. Kernel methods can be implemented in many supervised learning algorithms. However, not all learning algorithms support kernelization. For algorithms that do not support kernelization, basis expansion is a more efficient solution. The new features may help the learning algorithm to perform better in the transformed space than in the original space.

4.8. Model Selection and Evaluation

Identifying the correct machine learning algorithm and evaluating its performance are critical elements in a machine learning forecasting case. Frequently, no single model can give the user acceptable performance. Therefore, we need model selection and model evaluation issues to choose the most proficient model. Oftentimes, a single model can give an acceptable performance, and here, we need a single model choice. Therefore, we need a model selection issue to choose a specific model among the various modeling options the user resides in. In this section, we define some concepts and details regarding model selection and evaluation.

Types of Machine Learning Models

The machine learning algorithms currently available may be classified in various ways depending on the type of prediction task, by underlying structure, and by the specific supervised or unsupervised training mode.

There is no agreement on the machine learning methodology classification to be used. Throughout this text, we segregate methodology tests into two large classes: (1) supervised learning algorithms tackling both regression and classification problems and (2) unsupervised learning algorithms that can handle either clustering or dimensionality reduction. Finally, we can mention some methods that make both clustering and dimensionality reduction, embeddings, such as self-organizing maps, autoencoders, and generative adversarial nets.

Evaluation Metrics for Predictive Models

When evaluating predictive models, it is tempting to assume that any prediction error as a distance to a real value measured by any norm for a reasonably large value is good enough. However, distance has many interpretations, and any distance-error interpretation represents a systematic bias of the faulty model. Bias is a systematic prediction error, which is related to the distance function and the predictors distribution and is independent of the predictive model. The important point is to use a metric that is coherent with the user objective, thus avoiding a metric that would produce a systematic bias in the evaluation process. In the following list, we present what we consider to be commonly used evaluation metrics to model predictive error or predictive capacity. They differ on the distance functions they involve.

4.8.1. Types of Machine Learning Models

There are many contributions to different academic domains such as econometrics and statistics that can be complementary to the analysis of economic relationships. However, probably one of the main differences between econometric/statistical models and ML models is the willingness of the latter to accept a loss of interpretability to reap the benefits of much more accurate predictions. A key idea in ML that leads to this difference with econometric/statistical models is the generalization view. In broad terms, all predictive models aim to learn some underlying target function that generates the response variable from the distribution of the explanatory variables, and one straightforward way to assess the ability of a model to learn this underlying function is to estimate how well it performs in the prediction task on new data. We can say that a model has good generalization performance if it predicts well on new data, if it exhibits low generalization error. The generalization perspective leads to the "no free lunch principle" in ML, which states that no supervised learning algorithm is better than any other when the average performance is considered across all possible distributions of the input data.

Supervised learning could be roughly divided between regression and classification tasks, which differ in terms of the characteristics of the response variable. In regression tasks, the response variable is continuous, whereas selection tasks view it as categorical. Such a difference in the nature of the response variable leads to different types of machine learning models, with different mathematical characteristics. Regression models can also be classified into two large groups: linear and nonlinear models. We can transform the objective of regression functions by specifying different error loss functions, although the most common is the squared error loss function used in least squares. Furthermore, if the choice of the structure of the regression function is made independent of the error loss function, regardless of the specifics of the distribution of the explanatory variables, estimation of the coefficients using least squares leads to different constraints on the regression function than estimating the regression using weighted quantiles.

4.8.2. Evaluation Metrics for Predictive Models

Predictive models are created using historical data to inform future customer behavior. However, it is impossible to know the actual customer behavior before it happens. Therefore, at the end of model development, various metrics are used to evaluate how well the model predicts. When working with regression predictions, one metric commonly used is predictive accuracy, measured with the coefficient of determination. However, this metric should be avoided in model evaluation if the focus is on the forecast horizon for some model, where variable A is equal to or lagged behind variable B. The reason is that these models typically have predictive accuracy worse than the base model with the same variable only lagged by one time unit. For predictive accuracy, predicted values could be used to compute isoquant distances with respect to an earlier stage for predictions. Note that evaluating an earlier model might have occurred too long ago if there were many future shocks. Because this metric does not distinguish forecast errors with different signs, other metrics such as mean absolute percentage error or mean absolute error could also be used to evaluate predictive performance.

A different perspective is taken when evaluating the predictive ability of original and revised models. The focus is on the future distance of a predicted quantile to some level of quantiles to evaluate how much the prediction is beyond some warning threshold and therefore relevant for decision-making processes in the future. Both cross-validation and rolling-origin forecasting performance evaluation help show possible reduced predictive accuracy of various models with regard to structural time series models designed to handle specific conditions, like trend and seasonal changes over time, changing seasonality, or outliers. Neither of the two evaluation types relies on balanced goodness-of-fit or forecast error measures that may produce false conclusions of no predictive loss.

4.9. Challenges in Implementation

As many of the chapters in this essay highlight, Artificial Intelligence and Machine Learning are being increasingly adopted by organizations to generate Customer Insights and predict Market Trends and Demand. To do this requires implementing AI and ML models at scale and with an ever-increasing criticality in terms of accuracy. However, challenges remain with these implementations. At its core, this is driven by an organization's ability to collect high volumes of high-quality data, and build models that bring the customer within the loop while also maintaining fairness, privacy, and transparency.

Privacy is one of the biggest challenges anticipated by businesses today. Customers worldwide are becoming increasingly concerned about how the personal data and, therefore, privacy of their data, is used and secured by businesses. This has become more pointedly highlighted during the recent pandemic, driven by the increase in the digital footprint of individuals, and a focus on content that at times is based on personal experiences and perspectives. From research, a recent study stated that two-thirds of consumers have reconsidered how their personal data is used, with many of them, especially younger consumers being less willing to share data. The younger demographic consumers have indicated a greater belief that digital advertising has become more intrusive during the pandemic, fueling their concerns over how their data is used by businesses and the targeted ads that are served to them. A related issue faced by businesses today is ensuring that models are developed that do not discriminate against a group or individual while at the same time also ensuring that accuracy is maintained. Balancing performance and fairness through the proper application of ML models is critical to understanding the opportunities and successes of AI today.

4.9.1. Data Privacy Concerns

Despite the fact that a wealth of data can be gathered about every single consumer and used as additional features in these models, many consumers are becoming fed up with the extent of data collection and surveillance that is taking place today. Individuals are becoming increasingly aware of the role of businesses in the mass collection and use of their personal data, and view this as an infringement on their personal privacy, particularly in light of recent high-profile data breaches. Consumers also exhibited knowledge and concern regarding the many negative uses of personal data by organizations, which include targeted advertising, personalized pricing, and identity theft, as well as the potential for unintentional negative consequences from organizations using their data, such as spam emails or the unintentional enabling of faulty algorithms.

While consumers felt that some form of data collection was necessary for consumer or public safety, security, or protection against disaster or loss, there was a need for companies to better manage consumer expectations about information use. Most significantly, consumers expected organizations to be open and honest about how their data is used, particularly in “making a profit from their personal digital footprints.” Increasingly, consumers are asking why companies that collect their data in return for seemingly free services or lower prices have a right to collect, exploit, and monetize that personal information. In many instances, consumers are unwilling to authorize automatic data sharing between products that facilitate unique service linkages. Consider the widespread adoption of ad-blocker plugins on web browsers, along with the popularity of the “Do Not Track” option.

4.9.2. Bias and Fairness in Models

Machine Learning has made amazing improvements in accuracy and the successful implementation of perception tasks such as visual recognition and natural language processing. Data-driven models for such tasks have, however, also shown some dilemmas for adoption in real-world scenarios. One would be the fairness concerns in regards to the demographic attributes associated with the perception data that are being made use of to build such models. These biases are often shown to transfer from the data to the training process and eventually into the prediction process of the model, thus need also to be addressed in model training procedures. In market forecasting, companies use previous assessments of their customers or customer groups based on specific demographic characteristics, or a customer-specific product acceptance model built on these assessments, to forecast what new competing products would appeal strongly to their customers or segments. If such perceptions are inducing a response bias, then those customers or segments exhibiting the beseech might be subject to a different level of service differentiation than those who do not. It would of course be desirable to control such differentiated rather diverse service levels a priori, since otherwise those subjected to the hierarchical service levels ex post would probably feel treated unjustly and thereby irritated, if not insulted. In other areas of marketing, demographic differences between customers are often also used to segment markets for targeting promotions or model advertising expenditures to foster demand for existing brands and products. In the scope of Big Data, various information about customers such as name, credit card number, e-mail address, bookmarks, gender, marital status, watch lists, purchase history, product ratings, and links on social networks can be exploited. However, it may also happen that these targeted promotions are actually misfiring, and models are thus biased due to some overfitting to specific customer segments with missing or improper links to market outcomes, as there could also be customers not reachable by the promotion, either they dislike the product or advertisement or are not present in the target area, and hence misallocation costs would occur.

4.10. Future Trends in Machine Learning

Business trends are accelerating interest in and deployment of machine learning within the enterprise. Often, the mention of machine learning is in the context of AI. While not all machine learning is AI, nearly all AI is based upon machine learning. AI gives the organization the ability to perform higher-level tasks on customers, manufacturing, maintenance, and many other areas, freeing up time and cost savings from core processes. Three key areas are key in ensuring that organizations are well-positioned to successfully utilize AI: continued advancement of algorithms, integration with Big Data technologies, and ease of use with ML tools. Algorithms are core to driving the accuracy that organizations need in applying AI. These algorithms enable deep learning, neural networks and other computationally intensive techniques that allow organizations to acquire and act upon insights regarding their customers and markets. These capabilities would not be available without algorithms that are suitable for solving the underlying statistical and machine-learning questions and the programs that implement them.

As organizations gather more and more data and recognize that different types of data are needed to identify and fulfill customer needs, machine learning with Big Data moves from potential and pilot models into production. Technologies allow machine learning to take advantage of distributed databases and in-memory databases to execute more complex algorithms on data at the speed of business. Machine learning moves from being data scientist focused and becomes accessible to the business analyst who knows the business but may not be conversant in machine learning. New cloud-based technologies allow smaller organizations to take advantage of machine learning, AI, and the enabling big data technologies to drive their own market differentiation. The path to vertical solutions within industries gets mapped with templates and best practices that guide organizations down the path of building their own unique capabilities.

4.10.1. Advancements in Algorithms

Recent developments in algorithm design have ushered in a new era of machine learning, reflecting a syncopation of theoretical breakthroughs with practical needs arising within the larger computer vision community. The impressive accuracy of deep convolutional networks sparked an avalanche of real-world applications of these models, with video and image tagging, as well as face detection and recognition, quickly becoming ubiquitous on most mobile platforms. The beauty of deep neural networks is their universal approximation properties, both for supervised and unsupervised objectives. Indeed, a growing scientific literature is emerging that shows that if the goal is simply to model the conditional distribution of the observed data or of a response variable given the data, then deep architectures are likely to outperform more empirical approaches. Indeed, for various tasks, the benefit we gain from the additional data is leveraged by an ensemble of deep networks, each with a different random initialization, class pooling, or pair of data augmentations. Further research in these areas is likely to continue to yield considerable improvements.

The resurgence of neural networks has occurred concurrently with breakthroughs in the related areas of generative models and reinforcement learning. In generative modeling, deep architectures are playing a central role in the learning and targeting distributions of observed variables. Since the celebrated success of deep belief networks that leveraged greedy layer-wise pre-training, deep networks have been successfully applied to the denoising auto-encoder task and generative adversarial networks have further redefined the state of the art in modeling according to distribution objectives. Finally, deep learning models are also enabling research in reinforcement learning to achieve high-stakes goals, such as winning games that were previously thought to be unattainable. Here, the interaction between model-based and model-free approaches, as well as the interplay with parametric models, could guide the next paradigm shift.

4.10.2. Integration with Big Data Technologies

Large volumes of market-related data are being generated every moment by various data sources like social networks, media publications, online rating and shopping portals, city councils, cultural organizations, and weather services. The generated data can be classified into structured data, semi-structured data, and unstructured data. This data is characterized by the three “v”s (volume, velocity, variety). It can be used for a more exact solution of different market forecasting and customer insight problems by using advanced machine learning algorithms focusing on specific tasks.

The integration of machine learning with big data technologies can allow for significant improvements in the accuracy of various ML models, making it possible for applications to process both high-volume traditional data gathered in a structured form, and also high-velocity high-variety non-traditional data in different forms such as images or text. It can also allow for faster learning by the methods and algorithms by significantly reducing their run time at the stage of model training. On the other hand, large data volumes can also push the development of newer, more accurate, and faster ML algorithms, extending the potential and capabilities of machine learning technologies. By doing so, big data technologies can provide additional huge benefits for various machine learning applications in market forecasting and customer insights, allowing users to address various business tasks at a much greater speed and with an improved quality using probability-based business analytical applications.

4.11. Case Studies

In this section, we highlight a selection of actual cases and examples where machine learning has been effectively applied to market forecasting. A sampling of the exciting implementations is presented, with a focus on companies that have achieved innovation and corporate success through the innovative application of machine learning. We look across the domains of market forecasting, credit scoring and scoring, and customer behavior analysis, with a focus on applications that exploit recent advances and successes in machine learning methods. A wide variety of sectors are covered, including finance, telecommunications, travel, retail, insurance, packaged goods, gaming, and technology.

Successful Implementations in Market Forecasting

ZestFinance has subordinated credit and risk analytics for US-based lenders for a while. What's unique about the company is its application of algorithms such as deep learning to building predictive functions used to aid large lenders in scoring consumers. The models work with hundreds of thousands of disparate data elements, including payment life on previous loans taken out and even non-financial influence, such as the randomness of their search behavior. Although the data passed on to the company for scoring might be obtained directly from bureaus, what actually distinguishes ZestFinance is the development and real-time application of machine learning methods that keep on infinitely learning about new data and relations over time.

4.11.1. Successful Implementations in Market Forecasting

While the aggregate demand forecasting has received considerable accounting research attention, the application of Machine Learning methods in constructing and improving aggregate demand and supply forecasts has received less interest, in part due to the lack of operational or firm-theoretical justification for such models. A basic Machine Learning model, the gradient boosting algorithm, often shows superior predictive capabilities relative to a large number of standard econometric forecast combination models in forecasting GDP and its aggregate components. It is also found that the economically meaningful prediction intervals can be constructed using quantile gradient boosting and that these intervals can be used for loss-averse decision making. Overall, there is a large number of forecast comparison studies which show that Machine Learning models can provide more accurate forecasts in a number of applications relative to traditional econometric methods, especially for leading and short-term forecasts. In particular, many find superior forecast accuracy using Machine Learning methods in predicting the following variables: GDP growth; recessions and turning points of the business cycle; exchange rates; consumption; no-show counts; sales; oil prices; inflation; industry output; unemployment; housing starts and sales; retail sales; U.S. federal budget surpluses; high dimensional factor models. These results show that, despite the lack of theoretical considerations for such models, Machine Learning model forecasts are superior in a large number of macroeconomic applications and tasks.

4.11.2. Innovative Uses in Credit Scoring

The reduction of adverse credit selection is the basis of successful credit banking. With the growth of non-banking financial enterprises, the need to electronically, automatically, or semi-automatically complete the assessment of credit requests has grown intensely. In this context, various credit scoring techniques have been proposed since the 1960s. Since then credit scoring systems have improved and new models have been developed. Access to better information and advances in cartographic, information and computing technologies have not only made it feasible to build ever more complex models but have also increased the number of applications.

With adverse selection in credit applications being a long-standing problem in unsecured lending, machine learning has the potential to provide solutions that enhance current credit decisions based mainly on traditional scoring models. These methods promise better predictions based on wider segments of the client base and consider a large number of features and thresholds and deeper insights into customers. These additional benefits are essential to lenders as they manage their business. With the advent of new tools for data analysis, machine learning models have passed the test of time and complexity. These algorithms are no longer the domain of privileged users, and their advanced capabilities for analyzing and tackling complex business problems can be accessed by a wider audience from different areas of business. Subsidizing this accessibility is the increasing availability of large sets of structured and unstructured data. The enhancement of digital processes has internalized and stored information about clients, transactions, and external events.

4.11.3. Customer Behavior Insights from Machine Learning

One important application of machine learning is really understanding what customers are looking for, what they find attractive in general, and what is likely to lead them into a buying mood. Understanding customers' needs has proven to be a major driver of corporate profitability and stock returns. The application of machine learning in this setting is simply an acceleration of what good marketers have done for a long time. Marketers traditionally relied on a combination of survey data to understand which customer characteristics led to specific preferences and parsing of aggregate purchase data to find the segment of the population that bought particular kinds of products.

One difficulty with this approach is that stated preferences often differ from actual buying behavior. A major advance of machine learning is that it does not rely on stated preferences; instead, its foundational technique is the recognition that a large proportion of behavioral differences within a population are likely to be systematically correlated with features that describe the population.

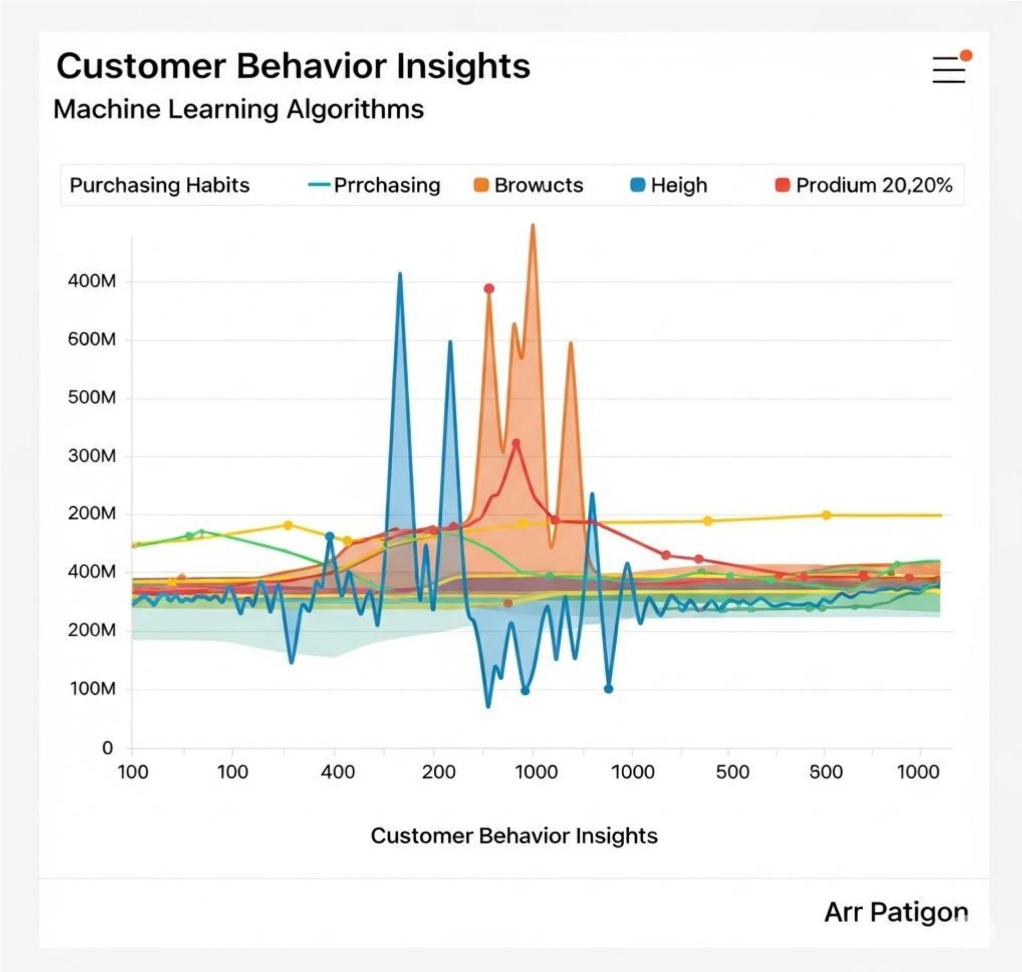

Fig : Customer Behavior Insights

Machine learning algorithms become fine-tuned to these population differences by drawing upon vast amounts of historical data on actual customer behaviors. The algorithms are then deployed to analyze real-time data streams to identify potentially key behaviors at any specific moment. They also identify possibly key customer characteristics correlated with these behaviors, particularly when the behaviors are compared with historic benchmarks.

4.12. Ethical Considerations

In recent years, there has been mounting debate regarding the usage of machine learning in important decision-making processes, including the use of algorithms for credit scoring, insurance underwriting, hiring, risk assessment, and other critical functions. Concerns include unintended bias in algorithms, lack of accountability when it comes to decisions made by machines, and lack of transparency in understanding how choices are made by machine learning algorithms. Bias in data used to train algorithms can lead to harm directed toward marginalized or under-represented groups. Even with fair training data, black-box algorithms produce biased results at a higher rate simply due to the way they make decisions, even when the output does not directly use sensitive variables. In the domain of finance and more specifically market forecasting and customer insights, we need to be particularly careful about ethical and responsible use of machine learning. When it comes to risk assessment, for example, historically biased data could be particularly pernicious and damaging. Unlike previous technologies, newly-minted machine learning models operate at scale and allow for instantaneous decision making. And unlike humans, machine learning models can be re-trained and re-tested in a short amount of time on different data, yielding a very different output almost instantly. This added scale and speed makes responsible AI use especially important for potential for harm. To that end, a number of principles have been proposed to help guide ethical and responsible AI usage, which we outline below. These principles have been re-applied and slightly modified from previous work.

4.12.1. Responsible AI Practices

Artificial intelligence (AI) is among the most transformative technologies of our lifetime, revolutionizing our personal and professional lives. The application of AI to consumer insights and market forecasting is growing at an exponential scale, and in many ways is in its infancy. Organizations have access to more data than ever to ground and power AI applications, most notably, the vast troves of unstructured data available through social media, search trends, earnings calls, and consumer reviews. AI is providing structure to this data, uncovering patterns in real time, to deliver insights that are actionable, explainable, transparent, and replicable in shorter and shorter time frames. Deep-dive essays are unlocked in a matter of minutes, and asynchronous and live translations are processed in seconds.

However, with this efficiency comes scrutiny. Given this transformation, many consumers – who increasingly demand companies "do the right thing" – are questioning how AI is being applied and whether it reflects our own ethical values and the governments we elect. Increasingly, business leaders are wrestling with questions of bias, fairness, accuracy, and privacy, and navigating the tension between moving fast and deploying cryptic algorithms that only a handful of experts can interpret, or creating accountable and explainable systems that are grounded on best practices. Data privacy was identified as the top factor for long-term global brand success. Likewise, a significant percentage of managers felt that their organization was operating in a manner that was not ethically responsible. As a result, companies must be wary of in which areas and how AI is applied, commit to not deploying it where it would conflict with ethical or moral standards, and to invest in Responsible AI best practices and programs to transform their organization.

4.12.2. Regulatory Compliance in Financial Services

Financial services are among the most regulated sectors of the economy. The regulatory landscape has created an environment in which the people behind regulated businesses and their technology suppliers closely monitor a number of regulations at the country and often local level. Broadly speaking, major areas for regulatory monitoring fall into the following groups: - Specialist financial services regulation: Regulations covering situations particular to financial services industry, for instance, can involve limitations around conflicted commerce, or solicitation of customers, denial of services to sectors of the economy such as "terrorist nations," mandates for assigning visible blaming for wrong-doing, and so on. - Anti-money-laundering (AML)/combatting-the-financing-of-terrorism (CFT): Regulatory regimes requiring banks and others to monitor potentially illegal transactions, which have gotten even stricter in the last decade. Inconsistencies and errors, which are traditionally not automatic, can incur severe penalties. - Data protection and privacy: Regulations that obligate service providers to protect identifiable customer data. GDPR-like regulations are increasingly being adopted throughout the world, with fines for failure to follow the letter or spirit of such regulations not uncommon. - Digital conduct regulation: A new type of regulation that has gained ground in Europe and other local jurisdictions in the past five years. Initially focused on technology companies, these regulations are growing their reach into finance, sometimes with strict limitations that can make transaction products uneconomical to provide, unless economies of scale can be achieved.

4.13. Conclusion

Machine Learning (ML) techniques are successfully introduced among many areas in business and science. In the present work, primary ML techniques were reviewed regarding new methods which were proposed for the Market Forecasting (MF) and Customer Insights (CI) problems. The aim was to provide a comprehensive description and to encourage both academics and practitioners to further develop for their use. As MF and CI problems are important but have never been treated in-depth, extensive attention has been paid to providing readers with substantial knowledge of these areas of research.

The conclusions we draw suggest that advanced ML techniques receive considerable success in introducing business/value innovation by proposing new methods. Technical and Market data become ubiquitous and it offers a new incremented knowledge of the customer. Usage of advanced ML techniques may provide automated intelligent systems that assist managers in their decision-making role. The next decade will see the appearance of a considerable number of scientific publications treating Market problems utilizing advanced ML techniques. However, scientists have a double responsibility. They should provide knowledge that is usable in business practice. And they also must be sure to consider ethical questions to avoid the possible introduction of biases during the resolution of specific Marketing problems.

Furthermore, Marketing could end up being a considerable source of knowledge for different areas in science as long as there is no used in excess and results are correctly documented to be usable. Of narrow question might attract attention among scientists in Finance as Market Forecasting will still be a vague area in decades or Internet of Things for Customer Insights Questions. In fact, Markets and Customers Intelligence will still provide an unclear area of investigation for years. The optimized usage of new and huge data in cyclic times will allow Machines Learning to help human beings in sound decision making so that cognitive biases are reduced.

References

I. Jain and K. Sharma, "Predictive analytics for market trends using AI," IJERU, 2024.

Y. Zhang et al., "Financial Customer Behavior Prediction Based on Machine Learning," ITM Web Conf., 2025.

M. A. Khan and S. Singh, "Credit Risk Prediction Using Machine Learning and Deep Learning," MDPI Algorithms, 2023.

R. Das, "The Role of Predictive Analytics in Enhancing Financial Decision," SCIRP, 2024. H. Li et al., "Big Data and Machine Learning-Based Risk Monitoring System," arXiv:2407.01562, 2024.