ISSN: 2959-6386 (Online), Vol. 2, Issue 1

Journal of Knowledge Learning and Science Technology

journal homepage: https://jklst.org/index.php/home

From Burden to Advantage: Leveraging AI/ML for Regulatory Reporting in US Banking

Sanjeev Prakash[1], Selvakumar Venkatasubbu[2], Bhargav Kumar Konidena[3]

[1]RBC Capital Markets, USA.

[2]New York Technology Partners, USA.

[3]StateFarm, USA.

Abstract

Machine Learning (ML) revolutionizes prediction processes, making them more cost-effective and precise. As the volume and diversity of financial data continue to grow, ML becomes increasingly valuable. One significant implication for regulators is the banking sector's growing reliance on ML methods for decision-making, which inherently lack full understanding by their creators. Consequently, regulators across all levels will increasingly encounter ML models that are challenging to fully grasp.Regulatory scrutiny is affected as supervisors must assess model risk. ML models incorporate numerous and intricate features, requiring examiners to comprehend their implications for transparency and associated operational risks. Moreover, utilizing historical data to train models may raise concerns related to fair lending practices. Already, some banks and FinTech firms employ ML across various banking services, including fraud detection, risk management, and pricing.Policy formulation may also feel the impact through two main channels: operational risk and market behavior. ML directly influences model risk, a subset of operational risk. Banks, bound by model risk management regulatory guidance established in April 2011, may find certain aspects of this guidance challenging to apply to ML tools due to their opaque nature. Furthermore, ML could potentially alter market behavior for certain liquid assets.

Keywords: Regulatory reporting, United States banking, Artificial Intelligence, Machine Learning, Automation, Compliance, Efficiency, Operational effectiveness.

Article Information

Article history: 13/01/2023 *Accepted:* 15/01/2023 *Online:* 30/01/2023 *Published*: 30/01/2023

DOI: https://doi.org/10.60087/jklst.vol1.n1.P176

Correspondence author: Sanjeev Prakash

Introduction

1. Overview of Machine Learning

Machine Learning (ML) represents an aspect of artificial intelligence (AI) that leverages computer systems to automate predictive tasks. While interest in AI has existed for decades, the fascination with ML has emerged more recently. Over the past decade, Google searches for "Machine Learning" have increased significantly, now surpassing those for its broader counterpart, "Artificial Intelligence," by double the volume. ML permeates various aspects of modern life, from suggesting new connections and products to facilitating autonomous driving. Its widespread applications have become so prevalent that educational tools now exist to aid children in comprehending their functionalities.

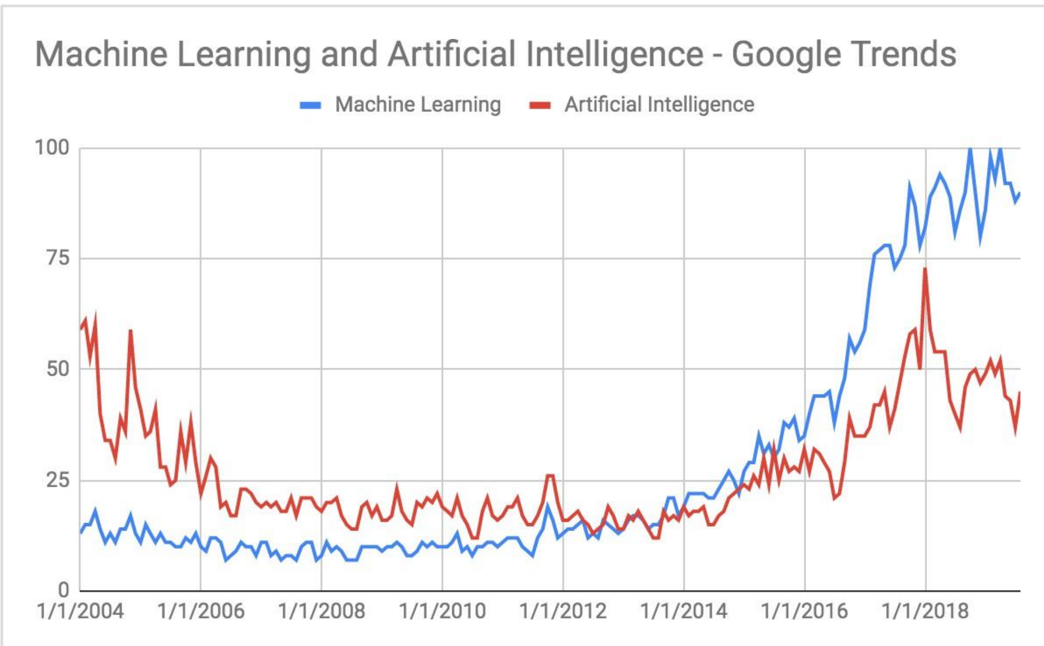

ML applications have historically been spearheaded by platform giants like Google and Facebook, but banks and Financial Technology (FinTech) firms are swiftly catching up. In 2017, Capital One established a dedicated "Center for Machine Learning" (C4ML) to support the development of ML systems across various business lines. Today, many banks integrate chatbots equipped with natural language processing (NLP) technology to enhance customer interactions and information retrieval. Moreover, models have long been integral to banking functions such as underwriting and fraud detection, and ML holds the promise of enhancing these existing models. Consequently, the increasing adoption of ML by banks comes as no surprise.

The incorporation of ML in the banking sector offers numerous advantages. ML streamlines prediction processes, potentially reducing costs for both banks and consumers. However, the adoption of ML also presents new challenges. ML methods automate prediction in ways that may diminish transparency, thereby increasing risks if not adequately understood and managed.

In this paper, we delve into the implications of ML for banking regulation, with many of these implications extending to the broader finance industry. Before delving into these implications, we provide an overview of ML to provide context for our discussion.

Why is Machine Learning (ML) a Hot Topic?

The resurgence of interest in machine learning (ML) can largely be attributed to the breakthroughs in Deep Learning. Initially developed in 2006, Deep Learning began to outperform other ML techniques by 2010, particularly when provided with ample data. Much of the notable progress in ML since then has been driven by incremental advancements in Deep Learning. As the volume of data generated and computational power continues to increase while becoming more affordable, the utility of Deep Learning and other ML methods is expected to grow correspondingly.

Another significant milestone was the application of Reinforcement Learning, which enabled AlphaGo to defeat the world champion Go player in 2016. Unlike Deep Learning, Reinforcement Learning can be effective even with limited training data, making it valuable in scenarios where machines learn through trial and error, such as in board games like Go.

Despite their impressive capabilities, both Deep Learning and Reinforcement Learning fall under the category of Artificial Narrow Intelligence (ANI). This means that these systems lack the ability to generalize their learnings beyond the specific tasks they were trained on. For instance, a Deep Neural Network trained to predict mortgage default rates cannot apply this knowledge to commercial loans or discern when to utilize a different model for subprime loans. Essentially, the underlying algorithms merely identify correlations between numerical input data and outcomes without grasping the broader context or underlying principles.

Where is Machine Learning (ML) Most Effective?

Machine Learning (ML) methods excel in addressing high-dimensional and intricate (non-linear) problems. There isn't a clear-cut threshold for determining when to employ ML, but it generally arises when imposing a predefined theory or structure on a prediction problem becomes overly limiting. In scenarios where prediction tasks are exceedingly complex, manually testing all potentially relevant relationships becomes impractical. In such cases, we depart from the assumptions of traditional statistics (TS) and allow the data to guide us. This approach often results in the development of black box models that have the potential to enhance predictive performance.

The ongoing digitization of data is expanding the scope of problems for which ML methods can deliver superior solutions. Consequently, both bankers and regulators will increasingly encounter these black box models or risk falling behind in the artificial intelligence race. Existing regulations crafted during an era of explicit underwriting and risk management systems, such as Fair Lending, Know Your Customer, and Model Risk Management, are already grappling with the challenges presented by ML methods, which inherently lack full comprehensibility due to their design.

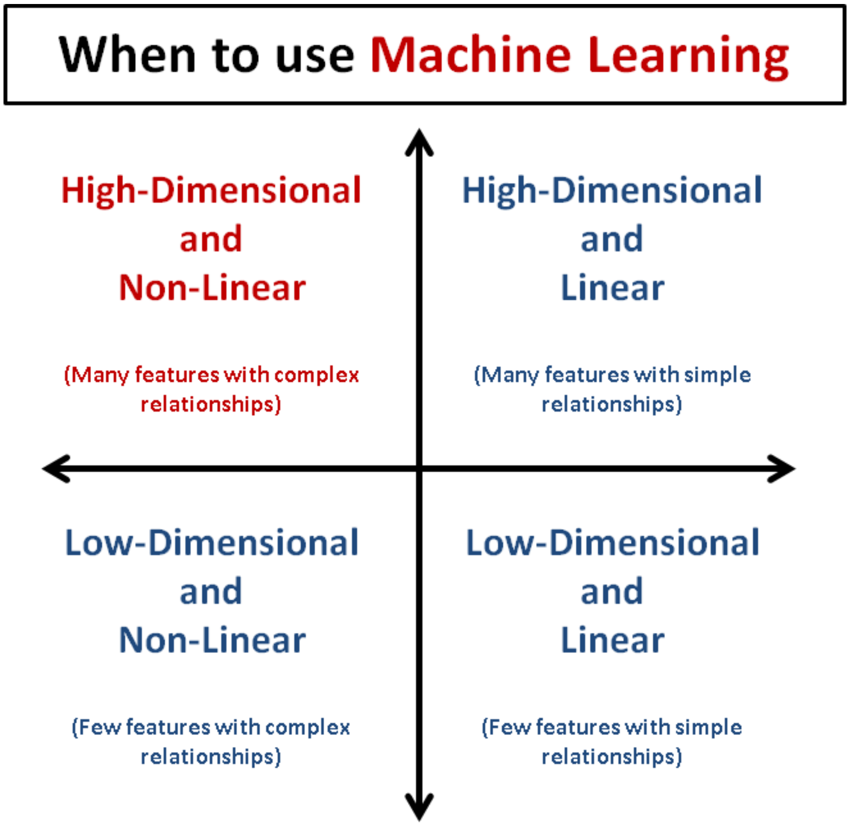

We've previously touched upon two crucial factors that delineate problems ideally suited for Machine Learning (ML): high dimensionality and non-linearity. Additionally, other factors such as domain expertise and the pace of regime change play significant roles.

Consider commercial credit modeling as an illustration. Proficient credit model developers are cognizant of the factors driving default prediction, often categorized as the "Five Cs": character, capacity, collateral, capital, and conditions. These seasoned modelers also understand the anticipated impact of these factors on default rates. Consequently, it might be prudent to confine the estimation process in alignment with the modeler's existing knowledge. This approach restricts the degrees of freedom to aspects that remain uncertain to the developer, such as the precise magnitude of the relationships.

Furthermore, the enduring relevance of the "Five Cs" suggests that they are unlikely to fade away anytime soon. Occasionally, however, the fundamental nature of a dynamic process undergoes a significant shift, termed "regime change". This phenomenon renders historical data and previously valuable domain knowledge obsolete. Integrating domain knowledge directly contradicts the essence of ML, which relies on data-driven insights. For instance, credit prediction models exemplify an area where regime change is likely to occur at a slower pace. The domain knowledge surrounding the timing and circumstances of borrower default is expected to remain relatively stable. While an unconstrained ML model may yield superior results with sufficient data, practitioners must weigh this improvement against the loss of transparency and potential emergence of non-intuitive relationships.

In contrast, traditional statistical methods continue to dominate in low-frequency scenarios such as macroeconomic forecasting and pricing illiquid assets like commercial credit. Although ML methods can offer assistance, the combination of extensive domain expertise and limited availability of relevant data complicates their application. This challenge is particularly pronounced in operations with less standardized data, such as commercial credit underwriting.

On the opposite end of the spectrum lies High-Frequency Trading (HFT), where ML-driven strategies defy human comprehension due to the rapid price fluctuations occurring at the nanosecond level and the vast volume of data involved. Additionally, in highly liquid markets, multiple ML-driven HFT strategies interact with each other, leading to swift regime changes, especially when the underlying estimation processes driving these strategies are dynamic. Unlike humans, dynamic HFT ML models automatically re-evaluate their trading strategies over short periods, such as daily or hourly. Picture multiple dynamic HFT ML models trading with and learning from each other in real-time—an endeavor beyond human capacity. Manually processing new data inputs, analyzing, re-estimating, and deploying new models into production slows down the adaptation process. Automating this process, however, has the potential to accelerate regime change significantly.

Origins of Big Data

In the finance sector, many machine learning (ML) applications necessitate access to "Big Data," a term often hyped to denote exceptionally large datasets. When we refer to "very large," we mean datasets that are so extensive that traditional models and analytical approaches become impractical. The volume of data generated from various sources such as mobile applications, payment systems, trading platforms, and even certain consumer credit products is staggering. Consequently, the industry increasingly relies on black-box and unsupervised methods to handle and derive insights from these massive datasets.

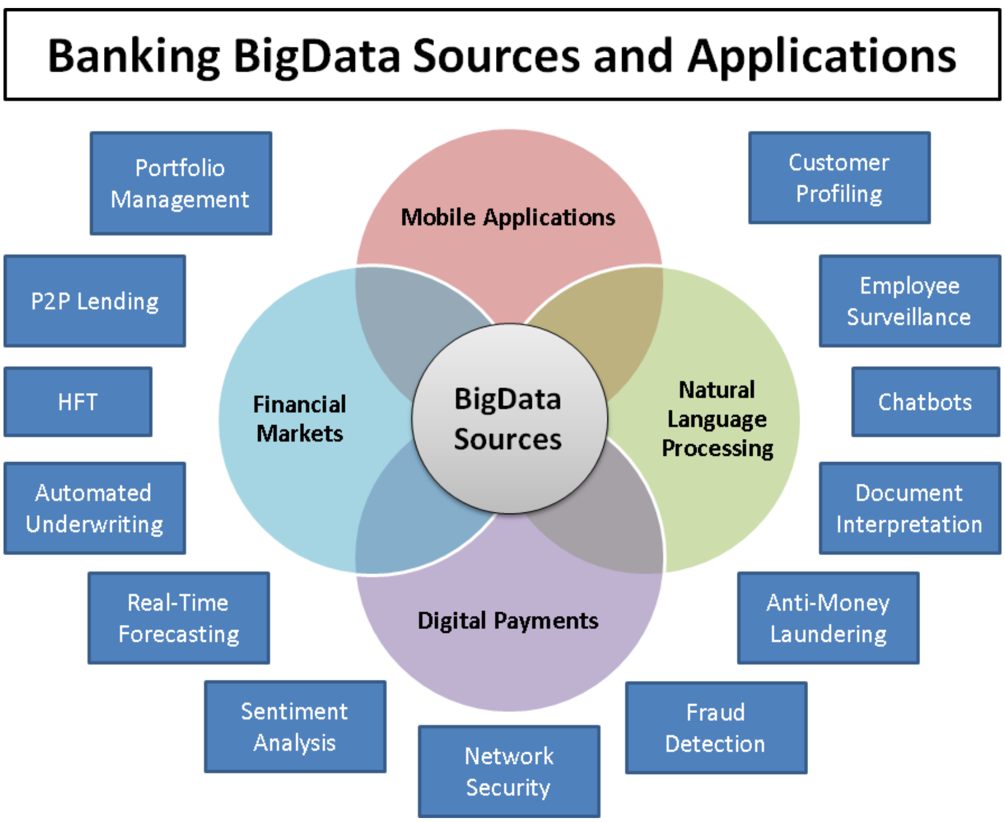

Understanding the origins of Big Data provides insight into the trajectory of machine learning (ML) applications. There are four primary sources of Big Data in the finance sector: Natural Language Processing (NLP), mobile applications, digital payments, and financial markets.

NLP, a subset of ML, presents novel opportunities for prediction by converting spoken and written language into actionable insights. Within the banking industry, NLP finds applications in information retrieval, intent parsing, sentiment analysis, speech recognition, and classification. For instance, information retrieval aids customers in finding relevant documents on a bank's website using keywords, while intent parsing involves understanding customer queries when interacting with chatbots and automated customer service applications. Banks are only beginning to tap into the potential of NLP to streamline operations, mitigate risks, and enhance customer services.

Mobile applications have revolutionized how consumers access financial services, concurrently generating new proprietary datasets tailored to specific services. When customers download mobile apps, firms gain access to valuable data such as customer locations, app usage patterns, and other mobile-generated data points that may not directly correlate with traditional labels of interest. For example, research has shown that seemingly unrelated factors like battery life can influence default probability.

Digital payments, including transactions between businesses (B2B), peers (P2P), and businesses and customers (B2C), are experiencing exponential growth. This proliferation of digital payment methods has led to a wealth of data on spending behaviors and cash flows previously obscured by cash transactions. In countries like China, the widespread adoption of mobile payment solutions like WeChat Pay and Ali Pay has virtually eliminated settlement times, bypassing traditional credit card usage. In the United States, companies such as Square, PayPal, Venmo, Apple, and Facebook are driving the adoption of digital payment methods, challenging traditional banks by offering a wide array of financial services.

Financial markets are rapidly transitioning to digital platforms, presenting a multitude of opportunities for machine learning (ML) applications in both traditional and emerging financial services. This digital transformation has paved the way for various ML applications, including peer-to-peer (P2P) lending, automated underwriting, high-frequency trading (HFT), and real-time asset pricing and forecasting.

Digitalization has facilitated the emergence of automated markets and streamlined price dissemination, offering fertile ground for ML-driven innovations. P2P lending platforms leverage ML algorithms to assess borrower creditworthiness and automate the underwriting process, while automated HFT strategies capitalize on real-time market data to execute trades at lightning speed. Additionally, ML techniques enable real-time asset pricing and forecasting, empowering financial institutions to make informed investment decisions in dynamic market conditions.

While there are numerous other sources of data contributing to the ML landscape, the primary drivers of ML adoption in the financial sector are Natural Language Processing (NLP), mobile applications, digital payments, and financial markets. These data sources represent the cornerstone of ML applications in banking and FinTech, shaping the trajectory of ML adoption and innovation in the industry. By understanding the dynamics of these data sources and the economics of ML, stakeholders can better assess the potential areas where ML will continue to proliferate.

Economics of machine learning (ML

The economics of machine learning (ML) are characterized by significant cost reduction in prediction tasks, primarily driven by automation. This understanding is pivotal in determining the extent to which ML will be adopted across various sectors. As costs decrease, financial institutions are increasingly inclined to deploy ML models in domains traditionally reliant on human judgment.

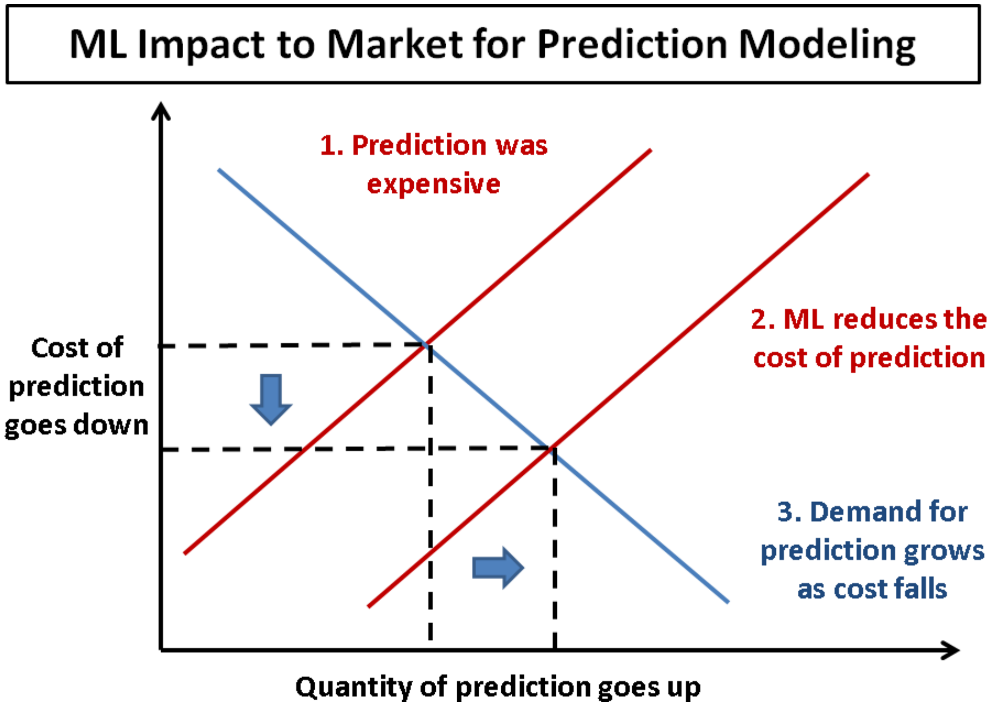

The advancement of machine learning (ML) has opened up entirely new avenues of inquiry, enabling automated prediction in areas where prediction was previously absent. This has significant implications, particularly for regulators who should anticipate ML methods impacting areas traditionally devoid of predictive models, such as customer interactions and information retrieval within banks. Regulators themselves may leverage models to assist in decision-making processes that were previously reliant solely on human judgment, such as reviewing credit agreements and examination reports.

This shift is occurring incrementally, with tasks previously unrelated to prediction now benefiting from basic ML tools. For instance, clustering algorithms can replace conventional methods of grouping publicly traded bonds or stocks by industry or market capitalization, instead utilizing hundreds or thousands of features for identification. Banks are already employing clustering algorithms to identify outliers in modeling data and detect potentially fraudulent transactions.

The economic principles of ML suggest that banks will continue integrating ML methods into various aspects of their operations where data is available. As discussed earlier regarding data sources, the volume and diversity of data are expected to expand. Banks embracing ML are likely to gain a competitive edge, driving the banking industry towards broader adoption of ML models and techniques.

Banking Applications of ML

Banks are harnessing the power of machine learning (ML) across various domains, including Natural Language Processing (NLP), risk and portfolio management, customer experience and behavior analysis, and fraud detection, including anti-money laundering (AML). While some applications are currently in use, others are theoretical concepts that may or may not be implemented at present. Due to the proprietary nature of many banking and FinTech models, our insights are limited, and we distinguish between existing and potential applications based on available information.

Natural Language Processing (NLP)

NLP is unlocking novel opportunities for prediction by converting spoken and written language into actionable data. In the banking sector, NLP finds applications in information retrieval, intent parsing, and classification, permeating various banking operations.

- Information Retrieval: Banks employ NLP to assist customers in finding documents on their websites and enhance employee performance. For instance, Capital One introduced an Alexa skill enabling users to check balances and pay bills via voice commands. JP Morgan Chase utilizes COIN, a ML software automating legal document review for credit contracts.
- Intent Parsing: Chatbots equipped with NLP capabilities aid in reducing reliance on human call centers, enhancing customer service, and recommending products. Kasisto's KAI chatbot assists bank customers in payments, transaction retrieval, and financial management.
- Classification: ML tools categorize standardized communications like emails and phone calls, enabling banks to flag anomalies. Some banks employ ML to detect suspicious employee activity, while regulators may leverage NLP to extract patterns from examination reports and financial statements.

NLP also extends its utility to risk and portfolio management, customer experience and behavior analysis, and fraud detection, as discussed further below.

Risk & Portfolio Management

Machine learning (ML) is exceptionally well-suited for numerous risk and portfolio management applications. From market risk management to asset pricing and algorithmic trading, ML methods offer enhanced accuracy and efficiency. Additionally, ML is increasingly applied in underwriting, credit analytics, macroeconomic forecasting, sentiment analysis, and document interpretation for mergers and acquisitions, enabled by natural language processing (NLP).

- Market Risk Management: ML methods can enhance market risk management by testing numerous potential market factors and accommodating complex relationships between these factors and portfolio values. While traditional methods rely on fixed sets of factors and linear assumptions, ML allows for the exploration of non-linear relationships, necessitating regulators to adapt to new factors and relationships while balancing accuracy gains with transparency concerns.

- Pricing and Trading Financial Assets: ML-driven electronic trading, even in less liquid markets like bonds, has surpassed $1 trillion in assets under management. ML models, which outperform traditional methods, leverage data acquisition and utilization, leading to the concept of "Algorithm Efficient Markets" (AEM). As automated ML continues to expand, the assumption that prices of liquid assets reflect all obtainable data becomes more plausible.

- Underwriting and Credit Analytics: ML advancements, including NLP, enable the approximation of credit scores using borrower online digital footprints. Big data in consumer credit markets facilitates ML application, although challenges such as fair lending requirements persist.

- Macroeconomic Forecasting: ML tools like principal component analysis (PCA) group similar macroeconomic variables into principal components, capturing shared underlying signals. New data sources, including Google and Twitter trends, enhance traditional macroeconomic forecasting by incorporating real-time sentiment analysis.

- Mergers and Acquisitions (M&A): ML, particularly NLP, automates parts of the M&A process, facilitating documentation interpretation and forecasting impacts on share prices. Goldman Sachs, for instance, employs models to forecast share price impacts in M&A scenarios, exemplifying the application of ML in this domain.

Customer Experience and Behavior

The aftermath of the Financial Crisis saw a surge in Fintech firms experimenting with innovative mobile applications, leading to a significant increase in the accessibility of banking services. This transformation has markedly enhanced user experience through the widespread adoption of chatbots, digital banking platforms, and personalized services.

- Automated Financial Assistants: Natural Language Processing (NLP) facilitates the development of automated financial assistants like Bank of America's "Erica," which discerns customer needs and tailors responses and actions accordingly based on behavioral patterns.

- Digital Banking: Fintech firms and banks are leveraging digital platforms to expand market share and enhance data inputs for ML models. Mobile and financial market data are utilized to automate underwriting, develop investment advising apps, streamline mortgage applications, enable peer-to-peer (P2P) lending, and support startup funding. The prevalence of digital wallets, exemplified by Tencent's "WeChat Pay" with 900 million monthly active users, underscores the decline of cash usage in favor of smartphone transactions, particularly evident in China.

- Personalized Customer Experience: Mobile app data facilitates the creation of detailed customer profiles, enabling tailored experiences and targeted marketing efforts. FinTech firms and banks leverage this data to customize services, allocate resources effectively, and market tailored offerings to each customer profile. For instance, Personics aided the Royal Bank of Canada in deploying a chatbot that learns from customer transaction patterns to provide money-saving recommendations. NLP-driven research and analytics, as demonstrated by SAS's analysis of Royal Bank of Scotland call center data, offer insights into customer complaints and preferences.

- Customer Lifetime Value Estimation: Enhanced data analytics enable more accurate estimation of customer lifetime value, allowing algorithms to recommend products, predict spending patterns, and refine banks' assessments of customer value.

Fraud Detection and Anti-Money Laundering

- Fraud Detection and Anti-Money Laundering (AML): ML is well-suited for fraud detection and AML due to the vast number of transactions and the rapid rate of regime change in criminal strategies. ML accelerates the detection process, crucial for minimizing losses in the ongoing arms race between fraudsters and regulators. Unsupervised ML methods such as clustering and classification are particularly effective for identifying suspicious anomalies. Given the dynamic nature of fraud, ML enables faster adaptation to evolving tactics and enhances the ability to detect new patterns of illicit activity.

- Data Security: ML contributes to enhancing data security by employing intelligent pattern analysis to identify sophisticated cyber-attacks. This involves a three-step process: clustering models identify patterns through unsupervised learning, experts evaluate these patterns to identify likely cyber-attacks, and ML models are trained using labeled data to detect future attacks in real-time, thus preventing security breaches.

Beyond these categories, various ML applications are emerging across diverse areas within the finance industry. While these applications may not neatly fit into predefined categories, they underscore the broad impact of ML. As the volume and diversity of data continue to expand, ML's influence in finance will continue to grow. However, regulators must carefully consider the trade-off between transparency and accuracy when implementing ML systems. Consequently, there are several implications of ML for banking regulators and financial regulators more broadly.

Implications for Banking Regulators

ML models, with their emphasis on accuracy over transparency and reliance on large datasets, present significant challenges for banking regulators. These challenges manifest in various areas, including model risk management, fair lending, transparency tools and techniques, new sources of fraud such as DeepFakes, feedback loops affecting market dynamics, and data economies of scale impacting governance and privacy considerations.

Model Risk Management:

- ML reduces human involvement in model development, potentially leading to less transparent outputs and increased model risk.
- Traditional model risk management guidelines may not fully accommodate the complexities of ML models, particularly in terms of feature engineering and black-box algorithms.

- Regulators may need to provide guidance on documenting and supporting ML models while balancing the need for transparency with model accuracy.

Fair Lending:

- ML methods in underwriting may raise concerns about fair lending practices, as models can inadvertently perpetuate biases or result in disparate impact on protected classes.
- Automated feature engineering in ML may introduce challenges in ensuring compliance with fair lending laws and regulations.

Transparency Tools and Techniques:

- ML's inherent complexity reduces transparency, posing challenges for model users, C-level leadership, and regulators in understanding model mechanics and outcomes.
- Regulators may need to adapt to the new vocabulary and methods associated with ML models to effectively assess risks and compliance.

New Sources of Fraud (e.g., DeepFakes):

- ML-powered banking services and applications increase the surface area for cyber attacks and fraudulent activities, necessitating enhanced fraud detection methods.
- DeepFakes present novel challenges in identity verification and fraud detection, requiring regulators to stay abreast of evolving technologies and their potential impacts.

Feedback Loops:

- ML models can influence market dynamics and regulatory outcomes, leading to both positive and negative feedback loops.
- Regulators need to anticipate and mitigate the effects of ML-driven feedback loops on market stability, risk management, and regulatory frameworks.

Data Economies of Scale, Governance, and Privacy:

- ML magnifies the advantages of data economies of scale, incentivizing banks to centralize and pool data.
- Regulators must balance banks' data incentives with privacy concerns and regulatory requirements, particularly regarding data governance and consumer protection.

Overall, banking regulators face significant challenges in adapting to the proliferation of ML models in financial institutions, requiring continuous monitoring, guidance, and regulatory frameworks to ensure stability, fairness, and transparency in the banking sector.

Conclusions

In conclusion, the increasing adoption of Machine Learning (ML) methods in the banking industry represents a significant shift in how financial institutions operate and deliver services. ML enables automation of prediction, making decision-making cheaper and more accurate, benefiting both financial institutions and their customers. However, this trend also poses challenges for regulators, as the use of ML often comes at the cost of transparency and raises concerns about fair lending, operational risks, fraud detection, and data privacy.

As ML applications continue to expand across various areas of banking, regulators must adapt to address these challenges effectively. This includes developing new approaches to model risk management, identifying and mitigating fair lending violations, monitoring the impact of feedback loops, and ensuring robust data governance and privacy protection measures.

Overall, while ML presents opportunities for improved efficiency and innovation in the banking industry, it also requires careful oversight and regulation to maintain transparency, fairness, and consumer protection. By staying vigilant and proactive in addressing the implications of ML, regulators can help foster a financial environment that leverages the benefits of technology while mitigating potential risks.

References List:

Abdi H., Williams L.J. (2010), "Principal Component Analysis," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 2, no. 4, pp. 433–459.

Adamopoulou E., Moussiades L. (2020), "An Overview Of Chatbot Technology," in: *Artificial Intelligence Applications and Innovations: 16th IFIP WG 12.5 International Conference, AIAI 2020, Neos Marmaras, Greece, Proceedings, Part II*, 16, pp. 373–383.

Aitken M., Leslie D., Ostmann F., Pratt J., Margetts H., Dorobantu C. (2022), "Common Regulatory Capacity for AI," [Online] Available: https://doi.org/10.5281/zenodo.6838946 [Accessed: 1.11.2022].

Ala'raj M., Abbod M.F., Majdalawieh M. (2021), "Modelling Customers Credit Card Behaviour Using Bidirectional LSTM Neural Networks," *Journal of Big Data*, vol. 8, no. 1, pp. 1–27.

Alzubi J., Nayyar A., Kumar A. (2018), "Machine Learning from Theory to Algorithms: An Overview," *Journal of Physics: Conference Series*, vol. 1142, [Online] Available: https://doi.org/10.1088/1742-6596/1142/1/01 2012

Artificial Intelligence: A Modern Approach, 4th Global ed. (2022), [Online] Available: https://aima.cs.berkeley.edu/global-index.html [Accessed: 11.02.2023].

Basu A., Manning W.G., Mullahy J. (2004), "Comparing Alternative Models: Log Vs Cox Proportional Hazard?," *Health Economics*, vol. 13, no. 8, pp. 749–765.

Bauwens L., Laurent S., Rombouts J.V. (2006), "Multivariate GARCH models: a Survey," *Journal of Applied Econometrics*, vol. 21, no. 1, pp. 79–109.

Biecek P. (2013), *Analiza danych z programem R. Modele liniowe z efektami stałymi, losowymi i mieszanymi*, Wydawnictwo Naukowe PWN, Warszawa.

Breiman L., Friedman J.H., Olshen R.A., Stone C.J. (1984), *Classification and Regression Trees*, The Wadsworth and Brooks, Belmond.

Buiten M.C. (2019), "Towards Intelligent Regulation of Artificial Intelligence," *European Journal of Risk Regulation*, vol. 10, no. 1, pp. 41–59.

Cath C. (2018), "Governing Artificial Intelligence: Ethical, Legal and Technical Opportunities and Challenges," *Philosophical Transactions of the Royal Society A*, vol. 376, 20180080.

Chan-Lau M.J.A. (2017), *Lasso Regressions and Forecasting Models in Applied Stress Testing*, International Monetary Fund, Washington.