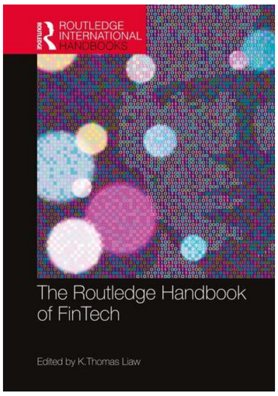

The Routledge Handbook of FinTech

K. Thomas

Machine Learning implications for banking regulation

Publication details
https://test.routledgehandbooks.com/doi/10.4324/9780429292903-29
Lihong L. McPhail, Joseph E. McPhail
Published online on: 15 Jun 2021

23

MACHINE LEARNING IMPLICATIONS FOR BANKING REGULATION

Lihong L. McPhail and Joseph E. McPhail[1]

1 Introduction to Machine Learning

Machine Learning (ML) is an application of artificial intelligence (AI) that uses computers to automate prediction. Humans have been interested in AI for many decades, but interest in ML is more recent. Google searches for "Machine Learning" have grown by an order of magnitude over the past ten years and now double those for its broader discipline "Artificial Intelligence". ML is everywhere; recommending new friends, products, and even driving cars. ML applications are so ubiquitous that tools are now available to help children understand how they work.[2]

ML applications have been led by platform companies like Google and Facebook, but banks and Financial Technology (FinTech) firms are not far behind. Capital One created a "Center for Machine Learning" (C4ML) in 2017 to help lines of business develop ML systems.[3] Many banks now employ chatbots that use natural language processing (NLP) technology to help retrieve information and improve customer experience. Models have already been at the heart of banking operations like underwriting and fraud detection for decades. ML has the potential to improve these models and so it should come as no surprise that banks are increasingly employing it.

Adoption of ML in the banking sector can provide many benefits. ML makes prediction cheaper and easier which may in turn reduce costs for both banks and consumers. However, ML also carries new challenges. ML methods automate prediction in ways that can reduce transparency. Reducing transparency can increase risks if not properly appreciated and understood.

In this chapter we discuss the implications of ML for banking regulation. Many of these implications have broader application to the finance industry. But before we discuss implications we first provide an overview of ML to help put these implications in context. In Section 2, we will cover common models and methods, including random forest, deep learning and artificial Neural Networks, and gradient boosting. In Section 3, we will discuss banking applications of ML, specifically natural language processing, risk and portfolio management, customer experience and behavior, and fraud detection and anti-money laundering. In Section 4, we will discuss implications for banking regulators, specifically model risk management, fair lending, transparency tools and techniques, new sources of fraud, feedback loops, data economy and scale, governance and privacy. The last section will conclude.

1.1 Why is everyone talking about ML?

Deep Learning[4] was the major breakthrough that re-ignited interest in Machine Learning (ML). Deep Learning was first developed in 2006,[5] but by 2010 started to surpass all other ML applications when given enough data. Much of what we see as rapid progress in ML is really incremental improvements in Deep Learning. As the amount of data we create and processing power continues to grow and become cheaper, so will the usefulness of Deep Learning and other ML methods.

Reinforcement Learning allowed AlphaGo[6] to beat the world's leading Go champion in 2016. Reinforcement Learning is distinct from Deep Learning[7] in that it is useful even when there is little data upon which to train. Reinforcement Learning is particularly useful in fields where the machine can learn from trial-and-error, such as a board game like "Go".

Both Deep Learning and Reinforcement Learning still fall into the category of Artificial Narrow Intelligence (ANI) which means the machine cannot generalize what is learned. Deep Neural Networks trained to predict mortgage default rates cannot generalize this knowledge to commercial loans, or know when to use a different model for subprime. The underlying code is still simply finding correlations between numerical input data and defaults.

Serious thinkers like Bill Gates and Sam Harris are concerned with the possibility of machines learning to generalize intelligence, a technology that is categorically different from what is available today called Artificial General Intelligence (AGI). While AGI has captured many imaginations and headlines, other serious thinkers consider AGI to be several major breakthroughs away. For example, Dr. Kai-Fu Lee noted, "There are no known engineering paths to evolve toward the general [AGI] capabilities." Implications of ML for regulators primarily come from the reduced transparency when implementing ML models such as Random Forests and Neural Networks, tools that are firmly within the scope of Artificial Narrow Intelligence.

1.2 What is ML?

ML is commonly defined as the field of study within AI that gives computers the ability to learn from data without being explicitly programmed.[8] A key source of confusion stems from the word "learning" which is analogous to "training" or "calibrating". All statistical models require training, but traditional statistics involves a lot of manual and often tedious steps. For this reason, we describe ML more simply as *any model that uses computers to automate prediction*.[9]

The term "automation" in the context of ML is a matter of degree. Likewise, the degree to which a model is "Machine Learning" is also a matter of degree dependent on the extent to which humans are manually involved in the prediction process. Model development, implementation, monitoring, and recalibration are all part of the prediction process. One or more of these parts can be automated. For example, a developer may judgmentally determine which features to include in the model, apply a Random Forest to structure how these features are mapped to a forecast, then recalibrate the model every quarter if forecasting errors breach some threshold. In this case, only a portion of the development process, the application of the Random Forest, is automated. The degree to which the developer manually sets "hyper-parameters", such as the number of trees in the forest, is another aspect of automation. A purer ML application of a Random Forest would automate everything from feature selection to forecast.

Experts sometimes disagree on which tools fall under the header of ML. For example, a common problem in macroeconomic forecasting is the need to combine many highly correlated features using Principal Component Analysis (PCA). James Stock and Mark Watson have been developing PCA methods using traditional statistics since Stock and Watson (2002). They didn't call their approach "Machine Learning" back then, but their approach is essentially an automated means of identifying relevant features and turning them into a prediction. Is PCA a type of ML? The answer seems to depend on who you ask. That's where it helps to think of ML in terms of automating prediction: If it automates prediction…then it is Machine Learning.

ML methods are not altogether different from traditional statistics (TS). Both ML and TS attempt to turn data into value, and use many of the same statistical building blocks. A key difference is in the applications most suited to ML methods (i.e. high-dimensional and complex problems).

Much confusion is caused by terminology. For example, ML practitioners refer to "labels" and "features" while TS practitioners use the terms "dependents" and "independents". Likewise with the ML terms "learning", "classifier", and "instance" which are analogous to the traditional terms "estimation", "hypothesis", and "data point". Practical applications of ML methods still require professionals to understand their objectives, choose the appropriate methods, compensate method strengths and weaknesses (such as over-fitting), and target the modeling solution to solve real business problems. ML is no magic bullet for bad data. *Garbage in* still results in *garbage out*.

ML models emphasize prediction accuracy over inference. Traditional statistical models are explicitly programmed in a manner that generically reduces to the form $y = \beta X + e$ where our dependent variable (y) is assumed to be a linear function of independent variables (X) with coefficients (β) and some error term (e). Explicit formulations help to infer the relationship between X and y.

ML and Big Data are often referred to in the same breath because *ML methods tend to outperform traditional statistics when applied to Big Data*. This is a direct consequence of relaxing assumptions (priors) about the relationship between a label and its potential features. Technically speaking, ML allows for far greater "degrees of freedom"; which simply means features and label interaction assumptions are relaxed to allow more potential solutions.

Today, there are so many ML methods that the term "Machine Learning" might seem difficult to define. Some ML methods, such as Reinforcement Learning, don't require much data. Others, such as Deep Learning, are more practical for problems involving Big Data. Random Forests and clustering models can work well with modest (traditional) amounts of data. What matters is that more prediction problems are being automated, making some harder to fully understand but allowing for potentially better predictions and new statistical tools, each with their own use cases, strengths and weaknesses.

1.3 Where is ML most useful?

ML methods are most useful for high-dimensional and complex (non-linear) problems.[10] No fine line exists to differentiate when to use ML, but practically it is the point where it becomes overly restrictive to impose a theory or structure on a prediction problem. Some prediction problems are too complex to manually test all potentially useful relationships. For these problems, we relax the assumptions of traditional statistics and let the data speak. The result is a black box that may improve performance.

Digitization is increasing the portion of problems for which ML methods can derive superior solutions. As a result, bankers and regulators will increasingly be forced to confront these black boxes or fall behind the artificial intelligence arms race. Regulations designed in an era of explicit underwriting and risk management systems such as Fair Lending, Know Your Customer, and Model Risk Management are already confronting challenges posed by ML methods that by design are not fully comprehensible.

We've already discussed two key factors that distinguish problems best suited for ML, namely high dimensionality and non-linearity (see Figure 23.1). Others factors include domain knowledge and pace of regime change.

Take commercial credit modeling as an example. Experienced credit model developers know the features that drive default prediction tend to fall into one of the "Five Cs" namely character, capacity, collateral, capital and conditions. The same modelers also know the direction in which these features should influence default rates. As a result, it may make sense to constrain the estimation process in a manner consistent with the modeler's priors. Doing so narrows the degrees of freedom to what the developer doesn't already know, such as the magnitude of the relationships.

Moreover, the "Five Cs" are unlikely to go away anytime soon. Sometimes the fundamental nature of a dynamic process changes. The technical term for this is "*regime change*". *Regime change* causes historical data and previously useful domain knowledge to become obsolete. Inserting domain knowledge goes directly against the grain of ML. Credit prediction models provide a useful example of where regime change is probably moving more slowly. Domain knowledge surrounding the nature of when and how a borrower will default is unlikely to change radically anytime soon. Given enough data, an unconstrained ML model may still produce superior results to constrained traditional statistics, but practitioners should weight this improvement against a loss to transparency and potentially non-intuitive relationships.

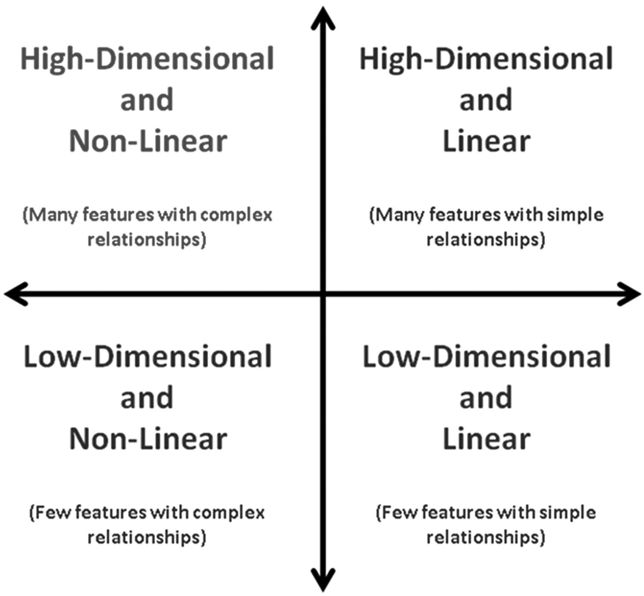

Figure 23.1 When to use Machine Learning
Source: Authors' analysis

Traditional statistical methods still dominate low frequency problems such as macroeconomic forecasting and illiquid asset pricing such as commercial credit. While ML methods may help the high degree of domain knowledge combined with low frequency of relevant data makes the application of ML methods more difficult. This is especially true for operations with less standardized data such as commercial credit underwriting.

High Frequency Trading (HFT) falls at the other extreme.[11] HFT strategies that are driven by ML are difficult for the human mind to comprehend due to the nanosecond price movements and large volume of data involved. Moreover, multiple HFT ML strategies are interacting with each other in some highly liquid markets. This can result in rapid regime change, especially when the underlying estimation processes driving these ML strategies are dynamic. A dynamic HFT ML model is automatically re-estimating its trading strategy over very short periods of time (ex. daily or hourly). Now imagine multiple dynamic HFT ML models trading with and learning from each other. Humans simply can't react to and learn from each other as fast as computers. We have to manually take in new data inputs, analyze, re-estimate, and put new models into production. Automating this process makes regime change potentially much more rapid.

We will discuss implications for dynamic ML models for banking regulation later, but one important consideration is the potential consequence of machines increasingly trading with and learning from other machines. Historical volatility, skew, return distributions, and other moments of price behavior might be changing faster today than before. Standardized measures of "riskiness" calibrated to the past decade or two might not accurately reflect their current risk and liquidity. Electronic trading is only a few decades old and has undergone major changes such as the rise of Exchange Traded Funds (ETFs), pricing to the penny (and smaller), HFT, ML, and an increasing proportion of trades performed between computers that learn from each other. Kearns and Nevmyvaka (2013) provide a nice overview of the application of Machine Learning in high frequency trading and market microstructure.

1.4 Sources of Big Data

Most ML applications in the finance industry need "Big Data", which is a hyped term that simply means very large datasets. By "very large" we mean so large that traditional models and analytics become impractical. The amount of data generated from some mobile applications, payment systems, trading platforms, and even some consumer credit products are incomprehensibly large, and hence the reliance on black-box and unsupervised methods (see Figure 23.2).

Understanding the sources of Big Data helps clarify where ML applications are headed. Four major data sources include Natural Language Processing (NLP), mobile applications, digital payments, and financial markets.

NLP is a field within ML that is opening up entirely new opportunities for prediction. By turning spoken and written words into features, NLP can turn records, text, and speech from customers (and employees) into valuable insights. Major applications of NLP in the banking industry include information retrieval, intent parsing, sentiment analysis, speech recognition and classification. Information retrieval occurs anytime customers are searching for information such as the use of keywords to find documents on a bank's website. Intent parsing includes the sometimes frustrating experience we have talking to chatbots and other automated customer service applications. As we discuss later, banks are just scratching the surface of opportunities to save time, reduce risk, and improve customer services via harnessing NLP.

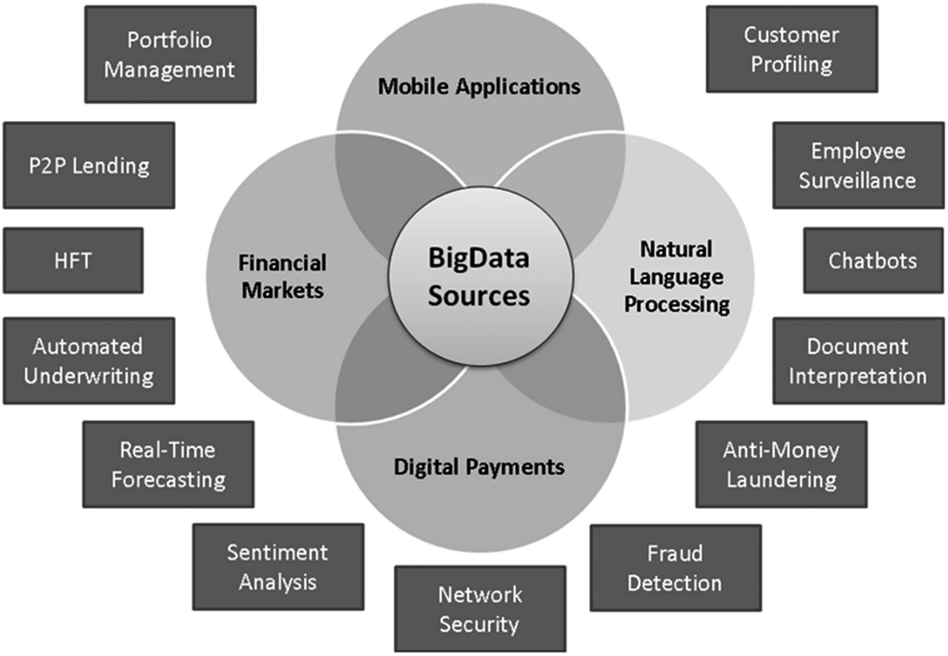

Figure 23.2 Banking Big Data source and applications
Source: Authors' creation

Mobile applications are revolutionizing the way consumers receive financial services. By creating apps, firms are also able to generate new proprietary data sources that are tailored to their particular services. Once a consumer downloads the app, firms might gain access to customer locations, apps they use, and other data generated from mobiles phone that just might have weak relationships with labels of interest. For example, battery life was found to be a weak feature in explaining default probability.[12]

Digital payments between businesses (B2B), peers (P2P) and between businesses with their customers (B2C) are increasing exponentially. This is creating a wealth of data on spending habits and cash flows previously obscured by cash. China has largely skipped over the use of credit cards and gone straight to mobile payment solutions like WeChat Pay or Ali Pay effectively reducing settlement time to zero.[13] US firms like Square, Paypal, Venmo, Apple, and Facebook are helping the USA to catch up in digital payments and competing with banks in a wide variety of other financial services.

Financial markets are largely and increasingly digital. Digital prices and automated markets have opened up many ML opportunities for traditional and new financial services such as P2P lending and other forms for automated underwriting, automated high frequency trading (HFT), and real-time asset pricing and forecasting.

There are certainly other sources of data, but these four compose the lion's share and are driving the areas where banks and FinTech firms are applying ML. Understanding these data sources and the economics of ML helps to assess where ML will spread.

1.5 *The economics of ML*

Automation is drastically reducing the cost of prediction. Understanding the economics behind this is key to assessing where ML will spread. Cheaper costs will lead financial

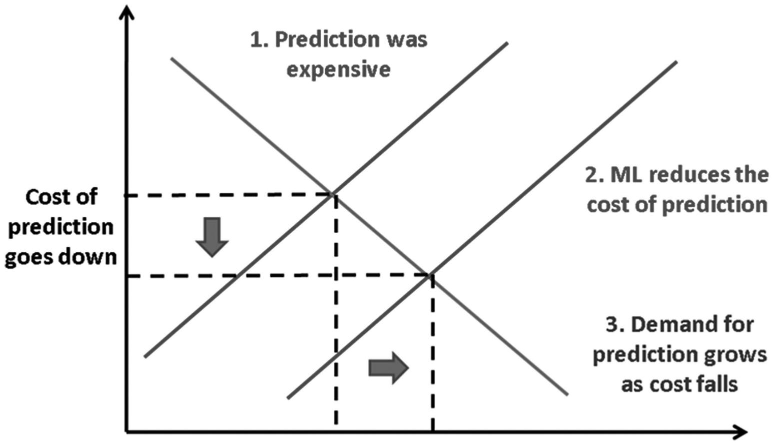

Figure 23.3 ML impact to market for prediction modelling
Source: Authors' creation

institutions to apply models in areas traditionally dependent on human judgment (see Figure 23.3).

Entirely new questions can now be asked. The result is automated prediction in areas where prediction did not previously exist. A key implication here is that regulators should expect more areas traditionally void of models to be impacted by ML methods such as how customers interact with the banks and how information is retrieved. Regulators themselves could employ models to help make decisions that previously relied entirely on human judgment such as the reviewing of credit agreements and examination reports.

This trend is happening "*at the margin*". Tasks that previously had little to do with prediction may now find benefits from simple ML tools. For example, instead of accepting traditional groupings of publicly traded bonds or stocks by "industry" or "market capitalization", clustering algorithms could be used to identify groupings based on hundreds or thousands of features instead of human judgment. Clustering algorithms are already being used by banks to detect modeling data outliers and potentially fraudulent transactions.

The economics of ML suggests that banks will continue to incorporate ML methods into every aspect of their business for which they have data. Our discussion on data sources suggests that the scope and magnitude of data will grow. Banks employing ML will likely have an advantage. As a result, competition appears to be driving the banking industry toward increased adoption of ML models and methods.

2 Common models and methods

ML models are often grouped into two distinct categories: supervised and unsupervised learning.[14] With supervised learning we know what the solution looks like, such as when predicting default rates, stock prices, or if a loan application is fraudulent. To predict defaults we explicitly assign a dependent variable (i.e. "label"). Three common supervised learning methods include Random Forests, Deep Neural Networks, and Gradient Boosting.

When labels are not given, this is known as "unsupervised learning"; such as when trying to cluster macroeconomic regimes or group companies based on similar characteristics. The most common unsupervised learning methods include clustering analysis and dimension reduction. We can use these tools to identify patterns and outliers. Banks employ "anomaly detection" to find potentially fraudulent transactions, but like all unsupervised methods, this ML tool doesn't have a clearly defined label like "Fraud vs Not Fraud". Humans need to come in after the fact to investigate if the anomaly is truly fraud or just an atypical transaction.

2.1 Random Forests

Random Forests are collections of many decision trees. Each tree uses an optimization algorithm to identify the feature with the most predictive power. This feature is called the "Root". Another algorithm separates the data at a point along the Root that best differentiates the label (see Figure 23.4).

For example, if we were predicting mortgage default rates, a likely root feature might be the borrower's credit score. This root is then split to optimally divide low vs high risk borrowers, perhaps at the score of 680. This process is repeated for subsequent (lower) features, such as the loan-to-value (LTV), until the algorithm runs out of features. The final node in the decision tree is the "Leaf" which contains the prediction. For example, a borrower with a FICO below 660, and a LTV above 90 might be found to have a 5% probability of default. In practice, Random Forests can be quite complex, consisting of hundreds or thousands of features, and pooling together predictions from potentially many thousands of individual decision trees.

2.2 Deep Learning and Artificial Neural Networks

Deep Learning is the ML method behind many developments in web search, advertising, computer vision, speech recognition, and self-driving. Deep Learning refers to training Artificial Neural Networks (ANN), a name that implies a crude approximation to a biological brain.

Brains contain billions of interconnected neurons which we use to learn through a process of repeated stimulation. ANNs are similar to their biological counterparts in that both rely on repeated stimulation to determine which interconnections to strengthen or weaken over time. However, that's about where the analogy stops, as biological brains are capable of general intelligence (discussed previously), a feat that no one has found a path for creating artificially.

Still, ANNs have captured many headlines and imaginations because their ability to solve clearly defined ("narrow") problems has grown exponentially. Deep Learning allowed ANNs to begin surpassing alternative ML methods such as Support Vector Machines in complex areas like image recognition around 2010.[15] Experts in ML don't fully agree on anything, but it seems clear to the authors that ANNs are winning in complex applications

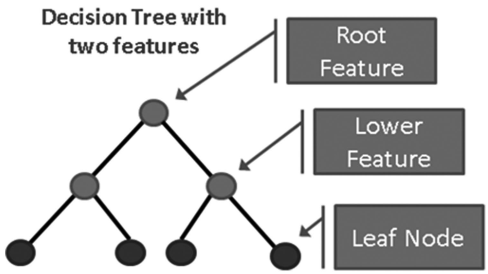

Figure 23.4 Simple decision tree
Source: Authors' creation

that have Big Data. Many big name companies use ANNs to drive these types of problems such as Facebook, Netflix, Twitter, Amazon, Tencent, Alibaba, Baidu.

The simplest ANN is called a perceptron, and is instructive in understanding the many types of more complex ANNs in use today. Perceptrons map input features to a binary output value. For example, given a borrower's FICO score and Loan to Value (LTV) on a new mortgage application, a perceptron can classify the borrower as either credit worthy, or likely to default. Figure 23.5 illustrates this process.

Perceptrons can have many inputs (features), each with its own weight which is optimized by training on real data in a manner that minimizes some error function. The greater the resulting weight, the more important the feature. The activation function f(x) takes the weighted sum, plus a bias to shift the activation function, and determines the resulting output.

By themselves, perceptrons provide little value. Today's ANNs have surpassed traditional methods (when given enough data) by allowing for more complex interactions between features. Deep Learning improved upon simple Neural Networks like perceptrons by including more complex interactions among features through "multi-layer" perceptrons (i.e. "hidden layers") and using "back-propagation" to effectively parse output error into improved neuron weights. This allows less important, but still useful, features to contribute to better forecasts.

There are dozens of ANNs today although the number continues to increase. Each has its own use cases. For example, Reinforcement Learning (RL) is particularly useful in static (unchanging) problems in which the range of possible actions can be clearly defined. This allows RL algorithms to generate data via trial and error, as did the Alpha Go machine previously discussed.

Recurrent Neural Networks (RNNs) are popular NNs used for time series forecasting or sequence data. RNNs retain information on previous states from one time period to the next by using forecasts as inputs in future periods.

2.3 Gradient Boosting

Gradient Boosting (GB) is a ML method that combines multiple weak models into a single prediction algorithm allowing developers to take advantage of approaches. Traditionally, developers would judgmentally select a single model, hopefully after verifying reasonable performance out of a sample under a variety of conditions. However, there may be particular

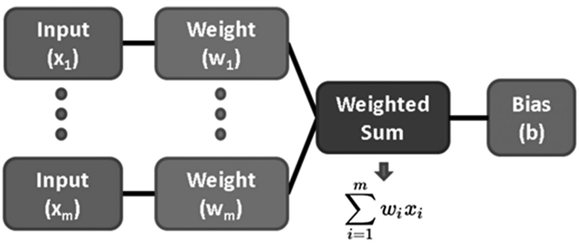

$$f(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{w} \cdot \mathbf{x} + b > 0 \\ 0 & \text{otherwise} \end{cases}$$

Figure 23.5 Perceptron (simplest ANN)
Source: Authors' creation

conditions (e.g. time periods, volatility regimes) where other models performed better. GB allows developers to automate this model selection process by sequentially adding multiple models trained on prior model residuals (i.e. "errors"). There is no limit to the number of additional models that can be applied in this way, although the greatest lift typically comes from applying models with varying strengths and weaknesses.

For example, suppose we already have a consumer credit card loss forecasting model estimated using a traditional reduced form regression formula. We have kept the formulation simple so that all features have a monotonic relationship to loss (our label). Now suppose we wish to employ a Random Forest in order to capture non-linear relationships that our original model is missing. To do this we would simply calculate the errors from the original model, and run through a separate Random Forest training exercise, being careful not to over fit to our training data. Combining both models would constitute a simple application of Gradient Boosting. We could automate this process further by letting the algorithm decide the sequence from a basket of models and model weights conditioned on performance.

Like Random Forests, Gradient Boosting is an "Ensemble" learning method because both combine multiple models to improve performance. However, gradient boosting is more general because it doesn't rely on a particular kind of model (such as decision trees).

Random Forests and Gradient Boosting lead many applications of ML within the finance industry. Both methods are quite powerful despite being much simpler and easier to implement than Deep Neural Networks.

What separates implementation of the supervised ML methods (like the three just discussed) from traditional methods is that they do not require the developer to explicitly state the relationships between input variables (i.e. "features") and the phenomenon being predicted (i.e. "label"). Traditionally, developers would need to painstakingly test residual patterns and interaction terms. Automating this process appropriately can save developers a lot of time.

Mmost powerful applications of Machine Learning in banking apply multiple models. For example, a model developer may apply principal component analysis to create new features, use anomaly detection algorithms to examine outliers, a Random Forest to identify and research the most useful features, and potentially Neural Networks when provided with enough data. ML model selection is less about *which is best* and more about *which is appropriate* for each problem.

3 Banking applications of ML

Banks are applying ML in a wide variety of ways. We discuss four broad categories including Natural Language Processing (NLP), risk and portfolio management, customer experience and behavior, and fraud detection including anti-money laundering. Some of the applications we discuss are currently being applied while others are theoretical applications that may or may not be in use today. We distinguish between the two based on our own observations which are limited due to the proprietary nature of many models used in the banking and FinTech industry.

3.1 Natural Language Processing

NLP is allowing entirely new opportunities for prediction by turning spoken and written words into readable data. Major applications of NLP in the banking industry include information retrieval, intent parsing, and classification. These applications are so broad that they permeate many other banking operations we discuss later in Section 5.

- **Information Retrieval** occurs anytime we search for information, such as a Google search. Banks are using information retrieval to help customers find documents on their websites, and help improve the performance of employees. For example, Capital One recently added a new "skill" to Alexa that allows users to verbally check their balance and pay bills. JP Morgan Chase launched a ML software called COIN, short for Contract Intelligence, which automates legal document review such as credit contracts.
- **Intent parsing** includes chatbots and other tools that discern context within human language. Banks are using chatbots for a wide range of applications such as reducing reliance on human call centers, aiding employees, improving customer service and recommending products. For example, Kasisto created a chatbot called KAI to help bank customers make payments, fetch transactions and account details (information retrieval), and manage finances.
- **Classification** includes tools for categorizing standardized communications like emails, chat windows, and even phone calls. Banks can use these categories to flag anomalies. For example, some banks are using ML to detect suspicious employee activity. Regulators may find that examination reports and financial statements can be mined for useful patterns in language that proceeded from failures in controls or relaxing of credit standards.

NLP is also being applied to the next three areas we discuss: risk and portfolio management, customer experience and behavior, and fraud detection including anti-money laundering.

3.2 Risk and portfolio management

ML is ideal for many risk and portfolio management applications. Market risk management, asset pricing, and algorithmic trading are all ideal because of the emphasis on accuracy over inference. ML is also being applied to underwriting and credit analytics, macroeconomic forecasting, sentiment analysis, and document interpretation for aiding mergers and acquisitions made possible by NLP.

- **Market risk management** may be enhanced by ML methods in liquid markets for which there is enough data. Traditional methods for measuring market risk tend to rely on a fixed set of "factors" such as a portfolio sensitivity to interest rates, inflation, a broad equity index such as the S&P 500, and commodities like crude oil. Relationships between a portfolio and these factors are typically assumed to be linear, meaning that a 1% change in a factor would lead to the same X% change in portfolio value regardless of the factor level. Both of these assumptions amount to an imposition of priors (assumptions) regarding how a portfolio's value can change in response to these factors. ML methods can be used to test hundreds or thousands of potential market factors and allow for more complex relationships between these factors and portfolio values. Regulators are accustomed to thinking about risk management in terms of these traditional factors, so incorporating new factors and non-linear relationships could take some getting used to. Market risk management is also an area where gains from accuracy clearly need to be balanced by any potential loss of transparency. ML models have been applied in all areas of banking risk management, such as credit risk, liquidity risk, and operational risk. Leo et al. (2019) provide a comprehensive overview of these applications.
- **Pricing and trading financial assets** – Investing is an arms race in which ML has long been applied. Electronic trading is growing steadily, even in less liquid markets

like bonds. [16] Automated trading has surpassed $1 trillion in assets under management, a significant portion of which is driven by ML models. Those that are better able to acquire and use data are more likely to outperform. This has led the authors to coin the term "Algorithm Efficient Markets" (AEM) to refer to a new efficient market hypothesis (EMH). As the size and impact of automated ML continue to grow, so should the assumption that prices of liquid assets such as larger corporate bonds and US equities reflect all easily obtainable and measurable data. This goes beyond the "Weak Form" EMH which states that only data on past prices, and perhaps trading volumes are fully incorporated into prices. When data on revenue growth, price-to-earnings ratios, liquidity coverage, and other balance sheet data become ubiquitous then it seems plausible to assume that much of this data is going to be priced in.

- **Underwriting and credit analytics** – Credit scores are nothing new, but Berg et al. (2019) show that these scores can be reasonably approximated using borrower online digital footprints.[17] Some of this is a product of advances in NLP. Some is also due to the ever increasing amount of information that individuals are willing to put online. Most consumer credit markets already involve Big Data. There are millions of mortgages and credit cards. Each loan also contains hundreds of records through time for tracking payments. Individual borrowers tend to be more homogenous in terms of risk characteristics (ex. FICO scores, LTV ratios) than corporate obligors which can have drastically different business models. All this homogenous data lends itself well to ML. New data sources from mobile applications appear to be making credit models even better. However, some regulations make applying ML fairly a challenge, such as fair lending requirements. We discuss this more in the next section.
- **Macroeconomic forecasting** typically deals with data at lower frequencies such as monthly or quarterly observations. That can make some ML models, such as Neural Networks, impractical. But many ML tools are being applied, such as principal component analysis (PCA).[18] Relationships between macroeconomic factors can be very complex and subject to regime change. PCA is being employed to group multiple similar macroeconomic variables into their "principal component", a shared underlying signal. New macroeconomic data and data relating to economic strength and investor sentiment are being developed. For example, Google and Twitter trends are increasingly being referenced and tied to traditional labels of interest such as consumer and business sentiment.
- **Sentiment analysis** can be performed using higher frequency data not dependent on surveys such as Twitter and Google Trends. Google searches for the terms "Layoff", "Recession", "Systemic Risk" and "Bankruptcy" were rising even before the financial crisis heated up in later 2008. Higher frequency data relating to macroeconomic health and stress opens the door for ML methods that benefit from more data like Random Forests and Neural Networks. Tensorflow is already being used to identify related words using Neural Networks.[19] This same technology could be used to identify when people are getting more concerned about the labor market or a recession.
- **Mergers and acquisitions (**M&A) have traditionally been very labor intensive due to the idiosyncratic nature of each individual company. But that hasn't stopped Goldman Sachs from automating parts of their M&A process and using models to forecast the impacts on share prices.[20] NLP has been a major breakthrough for many applications that at their core are really just documentation interpretation. M&A is just one example.

3.3 Customer experience and behavior

Since the Financial Crisis a wave of FinTech firms have begun experimenting with new mobile applications. As a result, the percentage of people with access to banking services has grown dramatically. User experience has been enhanced through mobile technology as FinTech firms and traditional banks expand the use of chatbots, digital banking, and personalized experience.

- **Automated financial assistants** are one direct application of NLP. Bank of America created a bot named "Erica" to act as an automated financial assistant. NLP can be used to detect what customers want and tailor responses and actions based on patterns of behavior such as which services to recommend.
- **Digital banking** – FinTech firms and banks have been moving to digital as a means to expand market share and in some cases expand the amount of data used as inputs to ML models. Data from mobile phones and financial markets are being used as inputs to automate and/or improve underwriting models, create apps for investment advising and personalized financial planning, automate mortgage applications, [21] allow P2P lending, [22] and fund startups.[23] As of 2017 there were 21 million digital wallets in the United States, [24] a drop in the bucket compared to Tencent's "WeChat Pay" which has 900 million monthly active users.[25] Cash is rarely used in China today as more and more transactions are conducted through smart phones.
- **Personalized customer experience** – Data from mobile apps make it easier to create customer profiles. This allows for a more personalized experience and opens opportunities for more direct marketing. This allows FinTech firms and banks to tailor the experience, allocate appropriate resources, and market tailored service offerings to each customer's profile. For example, Personics helped the Royal Bank of Canada employ a chatbot on their mobile app that is designed to learn from customer transaction patterns in order to recommend ways to save money.[26] NLP can also be employed for research and analytics. For example, SAS used NLP data from Royal Bank of Scotland call centers to help determine the cause of customer complaints.
- **Customer lifetime value** estimation can be greatly enhanced by the same data enabling more personalized customer experience. Algorithms that can reliably recommend which products to recommend and predict spending patterns can also be used to improve bank estimates of customer value.
- **Robo advisors,** also known as digital advisors, provide automated, algorithm-driven financial service planning**.** Robo advisors are commonly used for investment management, and can also be combined with other financial services, such as banking. For example, Citizens bank's SpeciFi Save & Grow allows its bank customers to specify how much of their balance they want to keep in cash and how much they want to invest. Then SpeciFi from Citizens Investment Services offers automated investment portfolios to its bank customers.

3.4 Fraud detection and anti-money laundering

- **Fraud detection and anti-money laundering** efforts are ideal applications for ML because of the enormous number of transactions that occur and high rate of *regime change*. Fraud is an arms race. As bankers and regulators begin to understand and clamp down on a particular strategy, criminals develop new strategies. This makes *domain knowledge* subject to some scrutiny, as what may have been true ten years ago, may no longer be true today. ML helps to speed up the detection process which is critical for limiting losses. Unsupervised ML methods such as clustering and classification are particularly useful here because they can help to detect suspicious anomalies.

- **Data security** has been a growing concern for both individuals and banks. ML can help by using intelligent pattern analysis to identify increasingly sophisticated cyber-attacks. Applications can involve a three step process whereby clustering models are used to identify patterns (unsupervised learning), experts then evaluate the patterns to identify likely cyber-attacks (labeling), which are then used to train ML models (supervised learning) to determine, in real time, which future behaviors are attacks before a security breech occurs.

A wide variety of ML applications are popping up in areas that don't neatly fit into the above categories but hopefully they provide a flavor of the areas where ML is having the biggest impact. Applications of ML in the finance industry will continue to spread because the amount and variety of data collected continues to grow. Of critical importance to regulators is the potential reduction in transparency when applying ML in favor of accuracy. As a result, here are several ML implications for banking regulators, and perhaps financial regulators more broadly.

4 Implications for banking regulators

ML models tend to emphasize accuracy over transparency, improve with more data, and expand to applications undergoing digitization transformation. These realities are the forces behind the implications of ML for regulators. We discuss challenges for regulators arising in the areas of model risk management, fair lending, operational risk, new sources of fraud including DeepFakes, feedback loops with the potential to change liquid market dynamics, and data economies of scale which may impact data governance and privacy considerations.

4.1 Model risk management

Model risk is, "the potential for adverse consequences from decisions based on incorrect or misused model outputs and reports".[27] ML essentially removes some of the human thinking involved in model development. This could include everything from feature engineering to prediction. This may lead to outputs or predictions less transparent than those from the traditional model, and consequently lead to an increase in model risk.

ML has typically focused on model estimation, but feature engineering (making dependent and independent variables from raw data) may take up much of the model development process.[28] Kanter and Veeramachaneni (2015) developed an automated feature engineering (AFE) algorithm that, when run through their preferred ML model (Random Forest) was able to beat a majority of human forecasters (along with their ML methods). They did not look at the raw data, did not select which variables to include in the final model, and still beat a majority of competitors in a fraction of the time.

The OCC published model risk management (MRM) guidelines in 2011 (formally adopted by the FDIC in 2017).[29] The first step MRM and Internal Audit groups are expected to take is an "evaluation of conceptual soundness" including a review of "empirical evidence supporting the methods used and variables selected for the model", "judgment exercised in model design", and "consistency with published research and with sound industry practice".

There is no ambiguity with this guidance when applying traditional statistical methods with their explicit formulations and thoughtful selection of model variables. Regulators have good reason for expecting banks to evaluate and document "*conceptual soundness*", "*variable selection*", and exercise good "*judgment*". However, it may be difficult in some circumstances for developers to innovate with new ML models if they are required to show "*consistency with published research and sound industry practice*". ML models by their very nature reduce human "*judgment*" by relieving developers of some tedious manual processes such as "*variable selection*". It would be impractical for a developer to evaluate the "*conceptual soundness*" of all the relationships embedded within the nodes of a Neural Network or branches of a Random Forest.

This excerpt is just one of many that might seem to go against the grain of ML. Regulators are already being confronted with these issues. They might provide guidance outlining how banks can better document and support AFE and ML black-boxes. Alternatively, banks may get better at opening these black boxes and providing insights into which features are driving the model, which interaction terms are material, and the conditional direction of key relationships. One thing we could see going forward is an emphasis on outcomes analysis and ongoing monitoring over theoretical justifications.

4.2 Fair Lending

Many FinTech companies such as ZestFinance have emerged since the Financial Crisis that use ML methods for purposes of underwriting.[30] On the surface, ML methods would seem to provide a solution to fair lending concerns. After all, how can a machine (unless explicitly programmed) be racially motivated? However, banks have been slow to adopt ML methods for underwriting in part because of Fair Lending concerns.

Fair Lending laws and regulations make a distinction between "overt evidence" and "disparate impact".[31] An underwriting process can be void of any component (human or machine) that is racially motivated, but still end up being detrimental to a protected class of applicants. For example, suppose that a bank's underwriting model penalizes mortgages loans under $60,000 because these tend to have higher default rates than larger loans. If these loan applications are disproportionately made by a protected class, then this could very well qualify as having a "disparate impact".

The use of automated feature engineering could pose a challenge here as well. Letting an algorithm loose on raw data without supervision may be totally appropriate when looking for anomalies in transactions as part of a bank's fraud detection strategy, but the same strategies could lead to an "overt evidence" violation. For example, small geographic areas are not typically allowed in underwriting models because of redlining;[32] which is "a form of illegal discrimination in which a bank provides unequal access to credit, or unequal terms of credit on a prohibited basis based on the location of the applicant's residence".

4.3 Transparency tools and techniques

The Achilles Heel of Machine Learning when it comes to banks is transparency. By design, ML improves accuracy by dramatically increasing the potential complexity of relationships between features and labels. But reduced transparency increases the risk of operational failures.

There are many ways reduced transparency can increase operational risks. The most obvious way as we have already discussed is model risk management. However, the obscurity caused by ML black boxes may extend to model users, C-level leadership, and regulators.

Model users have become adept at understanding traditional statistical models, as the banking world applied traditional statistics to ever more decisions such as the move toward automated underwriting. ML models require a new vocabulary that may take time to integrate its way into banking vocabulary, such as "Supervised Learning" and "Automated Feature Engineering". Many of these terms are not new to statisticians, but as the amount of data available continues to expand, the implications of these terms will become more relevant to banking and financial services more generally.

C-level leadership teams may grapple with the challenge of understanding the models driving an increasing portion of critical business decisions. Models are at the center of almost every large banking operation. ML models will only add to these transparency challenges.

Take the simple example of explaining the risk profile of an asset. Traditional statistical models such as linear regression can identify key drivers of an asset by modeling its price movement against a set of features the model developer believes to be relevant such as interest rates, currency movements, inflation expectations, credit spreads and ratings, and other factors. The output includes "Beta Coefficients" and "Standard Errors" identifying the sign and significance of each model feature. Oftentimes linear models are segmented and contain interaction terms, which can make identification of direction and importance a bit tricky, but nothing too difficult. Now consider a ML model such as a Random Forest. You can output the average feature importance level based on node placement, but there is no "significance level". There is also no "sign" to provide an easy check on the direction each feature has on the label. ML models are also likely to contain many more features than a traditional model. Now compound that problem across a growing number of bank models and exponentially growing amount of data and you start to see how transparency at the C-level may get more challenging.

Regulators may face many of the same challenges as model users and C-level bank executives. Regulators rely heavily on detailed documentation to identify areas of weaknesses in banking models and associated risks. ML models may not lend themselves as well to detailed documentation of model mechanics because the inputs and outputs are so much larger and complex. It is fairly easy to list out linear regression features and how these features were transformed before including them in a model. It may be impractical, if not impossible to do this for a multi-layer Neural Network with even modest numbers of features because of the compounding number of feature interactions. Stress testing outcomes and Value-At-Risk (VaR) measures may also become harder to compute.[33]

4.4 New sources of fraud (e.g. Deepfakes)

Banks and FinTech are making it easier for customers to access services by utilizing existing mobile applications and creating new ones. This adds value to customers, but may also be increasingly the surface area of potential cyber attacks and fraudulent behavior. For example, payment platforms like Square Cash allow users to create accounts automatically using very little customer information.[34] This makes it easier for customers to access traditional banking services like payments, but also relies heavily on ML methods to detect fake accounts and fraudulent transactions. Banks have been losing market share to FinTech firms since the financial crisis as customers have come to demand services that don't require coming into a branch and digging up documentation to verify identify and income. As a result, banks are competing by providing similar services.

Deepfakes provide a particularly unnerving example of how banks' arms race to apply ML to stay competitive could get tricky for regulators.[35] Academics first created the ability to create videos of people doing and saying things they have never done before using Machine Learning. When given enough video of actual people, Deepfake Neural Networks allow creators to replace videos of one person with another. The threat of this capability is pretty obvious. What is surprising is that to date we are not aware of any high profile fraud cases using this technology to impersonate another person. All it takes is video, and video sharing via social platforms has been growing exponentially.

Deepfakes today are not 100% persuasive, but they are getting better so fast that humans, at some point, might not be able to tell the difference. As a result, humans might need to rely on other ML algorithms to tell if a video of someone is real or fake. Clearly, we want the good guys to have access to as good or better technology than the bad guys. For this reason, the fight against fraud might be an area where regulators may allow unencumbered AI, even if that means less transparency into predictions and associated strategies by banks. However, unencumbered AI may have unknown consequences. Ongoing monitoring of innovative solutions could help identify these consequences before they have materially adverse impacts.

4.5 Feedback loops

What happens when models change the behavior of what is being modeled? Models have been impacting financial market prices since their creation. For example, persistent mispricing in options was partially eliminated after the Black-Scholes option pricing model became widely adopted. ML models are likely to also be having an impact on markets in which they operate.

Feedback loops can impact regulatory models as well. Suppose regulators were to develop a model for identifying banks that are at high risk of going bankrupt due to poor risk management. If the model works then over time a back test will reveal the model to no longer be effective because regulators will move quickly to correct the identified poor risk management practices. In this case the feedback loop is "negative" because the use of the model decreases the presence of the external signal it is attempting to model.

Positive feedback loops are more interesting and potentially more problematic. Suppose that a trading algorithm identifies a signal that has historically been reliable at predicting stock market crashes. Suppose further that the growing field of algorithmic trading firms using Machine Learning pick up on this same signal. What is the likely outcome? Assuming they all use similar historical data to train their models it seems more likely some may simultaneously act on the same signal. In light of this it seems understandable that some blame the increased use of algorithmic trading, some powered by ML, with an increase in the frequency of flash crashes.

Positive feedback loops tend to make historical patterns persist, even if humans would prefer to let these patterns die. For example, historical racial and gender bias in underwriting may get picked up in the depths of a Neural Network trained on historical data. Google ran into this issue when training their language translation Neural Network. When attempting to translate phrases like "They are a doctor" into English the original Google Translate would replace "They" with "He" because the historical training data had more instances of male doctors than female.

Artificial intellect may make the same mistakes as human intellect, only faster. As a result, feedback loops may become magnified. Positive feedback loops could lead to more volatility. Negative feedback loops could lead to faster price discovery. Interactions between ML trading models could lead to entirely new price dynamics that may cause historically calibrated risk-based capital requirements to become outdated. Regulators may want to consider how ML combined with feedback loops might impact market risk and the accuracy of regulatory capital levels that are calibrated to historical periods.

4.6 Data economies of scale, governance and privacy

Data creates economies of scale. More data means better decision making. ML increases the economies of scale created by data. As a result, it seems likely that banks with more data will have an increasing advantage in all the ML applications discussed.

Banks, therefore, have an incentive to grow market share and centralize their data. For example, suppose that the top three banks each control 10% of the credit card industry. Now suppose that two banks merge bringing their combined market share to 20%, or 2x the amount of the third bank. How much of an advantage does this much data give a bank in pricing credit card risk?

Perhaps not very much for plain vanilla households with good credit and regular income. But for atypical borrowers a 2x advantage could be very material. Lex Fridman, professor of ML at MIT, points out that deep learning is not good at dealing with "edge cases".[36] Overfitting leads to poor predictions. Borrowers with atypical characteristics are the "edge cases" of the lending world and banks with more data may have an easier time utilizing leading ML methods like deep learning which benefit from more data.

The number of banks in the United States has been falling since regulations preventing interstate banking were relaxed in the 1980s.[37] Much of this trend can be explained by mergers and acquisitions. As a result, many banks are challenged by data fragmentation. ML provides an additional incentive for banks to pool their data.

To maximize performance a ML algorithm would ideally have full access to all of a bank's data. However, regulatory and privacy considerations make this challenging. Allowing some data fields such as an address into an underwriting model could lead to red-lining violations. Allowing full use of all data available through a mobile application such as battery life, time used, applications downloaded, vocabulary, images saved, contacts etc., might create privacy concerns. The rapid expansion of data types and volumes may make it difficult to find the appropriate balance between bank incentives to grow and pool data and potential violations and privacy concerns.

5 Conclusion

Machine Learning methods automate prediction, making prediction cheaper and more accurate. The amount and variety of financial data will continue, and so will applications and the usefulness of ML. A key implication for regulators is that the banking industry is likely to rely increasingly on ML methods for decisions that, by design, cannot be fully understood by their developers. As a result of this, regulators at all levels may increasingly confront ML models at both the examination and policy level.

Areas where ML appears to be most useful today include natural language processing, risk and portfolio management, customer experience and behavior, and fraud detection including anti-money laundering. What these applications have in common is that they allow the use of big data and can have very complex relationships. This list of areas will likely expand over time due to trends in digitization, growing usefulness and expansion of ML tools, and incentives by market participants to use the best methods available to get ahead. It is difficult to see how ML will change how banks operate and how financial services are delivered, but the source of data provides a clue. Four major data sources include Natural Language Processing (NLP), mobile applications, digital payments, and financial markets (ex. prices, volumes). Many of the ML methods used by banks and FinTech firms are dependent on this data to create value.

ML creates both opportunities and challenges for regulators. ML is making prediction cheaper, and that's good for financial institutions, regulators, and customers because it allows for potential improvement in decision making. However, ML also creates challenges because the improved accuracy of ML models tends to come at a cost to transparency. This reduction in transparency combined with bigger data sets and digitization of previously unstandardized information has many key implications for how regulators think about model risk management, identify and avoid fair lending violations, address potential operational risks, observe impacts from powerful feedback loops, confront new threats from fraud like Deepfakes, and address incentives to grow from data economies of scale that may require enhanced data governance and privacy protection.

Notes

1 The research for thischapter was conducted and written by Lihong McPhail and Joseph McPhail in their personal capacity and not in their official capacity as an employee of the Commodity Futures Trading Commission (CFTC) or Federal Deposit Insurance Corporation (FDIC). The analyses and conclusions expressed in this book are those of Dr. McPhail and her co-author Joseph McPhail and do not reflect the views of other employees of the CFTC Office the Chief Economist, other CFTC staff, the Commission itself, the FDIC, other FDIC staff, or the United States Government. For inquiries please email either lihong.l.mcphail@gmail.com or joseph.e.mcphail@gmail.com. Special thanks to Xiaojian Zhao PhD and Ilya Rahkovsky PhD for your insights. We are very appreciative of the many friends and colleagues that provided comments on earlier versions of the chapter.

2 The website Machine Learning for Kids (https://machinelearningforkids.co.uk/) builds on a simple coding language called "Scratch" to illustrate examples of ML for kids such as object recognition.

3 "Machine Learning in the Enterprise: Lessons from the Front Lines", by Dave Castillo https://www.capitalone.com/tech/machine-learning/machine-learning-in-the-enterprise-lessons-from-the-front-lines/

4 Deep Learning is the process used to train Deep Neural Networks, which are Neural Networks with more than one "hidden layer" separating features from predictions. Hidden layers are used to capture interactions between features.

5 https://beamandrew.github.io/deeplearning/2017/02/23/deep_learning_101_part1.html

6 "AlphaGo Zero: Starting from Scratch" by David Silver and Demis Hassabis, https://deepmind.com/blog/article/alphago-zero-starting-scratch

7 "Artificial Intelligence: What's The Difference Between Deep Learning And Reinforcement Learning?" by Bernard Marr, Forbes, Oct. 22, 2018. https://www.forbes.com/sites/bernardmarr/2018/10/22/artificial-intelligence-whats-the-difference-between-deep-learning-and-reinforcement-learning/#6b713cb6271e

8 This particular definition comes from Arthur Samuel (1959).

9 Our views on ML were strongly influenced by, "Prediction Machines: The Simple Economics of Artificial Intelligence," by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. We agree with the authors' view that the impact of ML is most accurately understood as "a drop in the cost of prediction".

10 We thank Xiaojian Zhao, Ph.D. in Computer Science and a thought leader on machine learning for sharing this theoretical framework for identifying problems most suitable for ML methods.

11 HFT is a subset of automated trading, a broad concept that captures much more than ML strategies. HFT includes any program trading platform that transacts a large number of orders in less than a second.

12 "AI Superpowers: A Conversation with Kai-Fu Lee," https://www.youtube.com/watch?v=oNAFI3Lh97Y
13 "China's Digital Economy: Opportunities and Risks" by Longmei Zhang and Sally Chen, IMF Working Paper, January 2019.
14 "Machine Learning Mastery with Python", ebook by Jason Brownlee, available at https://machinelearningmastery.com/machine-learning-with-python/
15 "Deep Big Simple Neural Nets Excel on Hand-writen Digital Recognition", by Ciressan et al. https://arxiv.org/pdf/1003.0358.pdf
16 "Electronic Trading in Fixed Income Markets", BIS, Jan. 2016. https://www.bis.org/publ/mktc07.pdf
17 Berg, Tobias and Burg, Valentin and Gombović, Ana and Puri, Manju, On the Rise of FinTechs – Credit Scoring Using Digital Footprints (July 15, 2019). Michael J. Brennan Irish Finance Working Paper Series Research Paper No. 18-12, Available at SSRN: https://ssrn.com/abstract=3163781 or http://dx.doi.org/10.2139/ssrn.3163781
18 As mentioned previously when discussing James Stock and Mark Watson (2002), PCA is one of many methods sometimes called "machine learning" that not all practitioners believe is strictly machine learning.
19 https://www.tensorflow.org/
20 "AI Revolution Disrupts Investment Banking," by KC Cheung, May 2020. https://algorithmxlab.com/blog/artificial-intelligence-revolution-disrupts-investment-banking/
21 *For example, Rocket mortgage https://www.rocketmortgage.com/*
22 For example, Lending club https://www.lendingclub.com/
23 For example, Kabbage https://www.kabbage.com/
24 "Will Bank Branches Give Way to Digital Wallets?" by Bhavana Yarasuri, ARK Invest Dec. 2018. https://ark-invest.com/articles/analyst-research/digital-wallets/
25 *"One photo shows that China is already in a cashless future," by Harris Jacobs, May 2018. https://www.businessinsider.com/alipay-wechat-pay-china-mobile-payments-street-vendors-musicians-2018-5*
26 https://personetics.com/
27 *"Guidance on Model Risk Management," SR 11-7, Board of Governors of the Federal Reserve System, April 4. 2011.*
28 Actual time spent on any particular stage in a model development process can vary widely. The point here is that many real world financial prediction problems require data that requires standardizing, cleaning, bucketing, splines, transformations, interactions, and other traditionally time consuming processes. In contrast, later development stages such as estimation, production, and monitoring tend to be quicker because of standardized programming packages.
29 "Supervisory Guidance on Model Risk Management," https://www.occ.treas.gov/news-issuances/bulletins/2011/bulletin-2011-12a.pdf
30 https://zest.ai/product
31 https://www.fdic.gov/resources/supervision-and-examinations/consumer-compliance-examination-manual/documents/4/iv-1-1.pdf
32 https://www.fdic.gov/resources/bankers/fair-lending/
33 Typically, an examiner could use a bank's stress testing model by identifying critical risk factors and stressing them to some confidence level such as two standard deviations from the mean. This exercise is not as straight forward with a Random Forest or Neural Network. First, there is the previously mentioned problem of not easily determining which features are "significant" and the "sign". Then there is the issue of conditioning these features on all the other features. The "significance" and "sign" can change when any of the other feature values change. This makes it harder to generate a stress test based on a traditional measure of stress such as a two standard deviation move in risk features.
34 https://cash.app
35 "It's Getting Harder to Spot a Deep Fake Video", Bloomberg Quicktake, https://www.youtube.com/watch?v=gLoI9hAX9dw
36 Lex Fridman's lecture on Artificial General Intelligence. https://aiatadams.files.wordpress.com/2018/01/lecture1.pdf
37 "Interstate Banking" by Daniel Liberto. https://www.investopedia.com/terms/i/interstate-banking.asp

References

Agrawal, Ajay, Gans, Joshua, & Goldfarb, Avi. 2018. "Prediction Machines: The Simple Economics of Artificial Intelligence." *Harvard Business Review Press*, April 17, 2018.

Berg, Tobias, Burg, Valentin, Gombović, Ana, & Puri, Manju. 2019. "On the Rise of FinTechs – Credit Scoring Using Digital Footprints." Michael J. Brennan Irish Finance Working Paper Series Research Paper No. 18-12. Available at SSRN: https://ssrn.com/abstract=3163781 or http://dx.doi.org/10.2139/ssrn.3163781

Brownlee, Jason. "Machine Learning Mastery With Python." Ebook available at https://machinelearningmastery.com/machine-learning-with-python/

Lee, Kai-fu. 2018. *AI Superpowers: China, Silicon Valley, and the New World Order*. Houghton Mifflin Harcourt Company, 1st edition.

Leo, M., Sharma, S., & Maddulety, K. 2019. "Machine Learning in Banking Risk Management: A Literature Review." *Risks*, 7, 29.

Kanter, Max, & Veeramachaneni, Kalyan. 2015. "Deep Feature Synthesis: Towards Automating Data Science Endeavors." IEEE.

Kearns, M., & Nevmyvaka, Y. 2013. "Machine Learning for Market Microstructure and High-Frequency Trading." *Chapter in High Frequency Trading – New Realities for Traders, Markets, and Regulators, eds* David Easley, Marcos López de Prado and Maureen O'Hara, Risk Books.

Samuel, Arthur L. 1959. "Some Studies in Machine Learning Using the Game of Checkers." IBM Journal of Research and Development, 44: 206–226.

Stock, James, & Watson, Mark. 2002. "Forecasting using Principal Component From a Large Number of Predictors." *Journal of the American Statistical Association*, 97(460): 1167–1179.