On Revolution and Evolution of ML

A historiographic review on the application of Machine Learning in Finance in the last decade

Carlotta Berto Salmerón
September 2021

Fundación Rafael del Pino
Princeton University
Global History Lab

Contents

List of Figures

1 Introduction

This report examines the current situation of Machine Learning applications to the financial environment and explores its yet unearthed potential through the voice of experts in the field.

It provides a technical and historical contextualization on self-learning algorithms and a critical discussion on the most challenging and relevant aspects associated with their use in today's industrial scene. The latter is mostly based on the interviews conducted with technical experts and professionals in the field Dr Igor Halperin, Dr Carlos Jaureguizar and Ms Sandra Nieto ([7][8][9]).

1.1 Motivation

Technological developments carry structural changes in the way society coexists, produces and evolves.

The extraordinary advances that have taken place in the last 20 years have led to what is now known as the Fourth Industrial Revolution, a more than ever fundamental change in the way society is set.[10] The overwhelming force and speed of this change has been making us reconsider not only fundamental aspects of society as we know it, such as the way value is created but even what truly means to be human.

In particular, the unprecedented evolution in generation and storage of information is giving rise to a novel way of understanding relations, processes and global systems. In the era of Data, the tools capable of controlling it are destined for success.

Machine Learning (ML), a subset of Artificial Intelligence, is one of these. Its use is becoming so widespread that it can be found in almost any sector and the results it delivers are decidedly promising, if not already transformative. However, the distinct characteristics of the tool, prevents Machine Learning algorithms from being applied equally across fields and situations. A particularly interesting setting, because of the peculiarity of the laws by which it is governed, is that of Finance.

1.2 Literature Review

1.2.1 What is Finance?

Finance is the (social) science around the management of money. Economics is defined as the study of how individuals and societies choose to produce, distribute, and consume scarce resources, while Finance is the specialized branch of economics concerned with the origination, management and study of what makes up financial systems, that being money, banking, credit, investments, assets, and liabilities among others. [1][11]

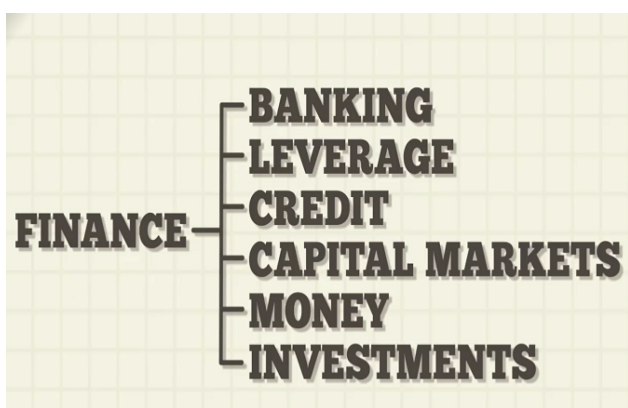

Figure 1: The many elements of finance.[1]

Many of the financial basic concepts come from micro and macroeconomics theory such as the *time value for money theory* [1].[1]

[1] A fundamental theory in macroeconomics, which essentially states that a dollar today is worth more than a dollar tomorrow.

Finance can be divided into three subcategories:

Public finance involves public systems such as taxation, government expenditure, debt and budget structuring and policy instruments (monetary for instance) among others.

Personal or private finance encloses the financial decisions and actions a household, as an economic unit, or a single individual partakes in. This might include mortgage planning, saving, personal investments and the like.

Corporate finance comprises the management of assets, revenues, debt and liabilities for a business. Some practical examples involve rates of return, cost of capital, risk quantification and optimisation of financial structures.[1]

In this report we will be mainly concerned with Corporate finance, since it is the area of finance where Machine Learning has been applied most.

1.2.2 What is Machine Learning?

Machine Learning (ML) is a subset of Artificial Intelligence.

By **Artificial Intelligence (AI)** it is meant *the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable*, as defined by one of the fathers of AI, John McCarthy. However, it is more broadly defined as the capability of a machine to imitate intelligent human behavior.[3]

Although the field of AI seems futuristic and even out a Sci-fi universe, it is nothing new. After WWII, some researchers already started to independently work on *intelligent machines*. The English mathematician Alan Turing may have been the first to do so and perhaps the first to decide that AI was best researched by programming computers rather than by building machines. [12] However, the field of AI itself was officially born in 1956 at Dartmouth College, where J. McCarthy coined the term for the first time. [13] By the late 1950s, there were many more researchers on AI [12] and by the mid 1960s laboratories were flourishing around the world.[14]

The ultimate goal of the AI movement was to make a machine intelligent enough for it to handle any general cognitive task presented in any setting, just like humans do. That is easier said than done.

Nowadays, the general AI ecosystem is classified into two branches: *weak* (or *narrow*) and *strong* (or *general*) AI. The distinction arises from the domain in which the AI system is able to "prove its intelligence" in. The more the system approaches characteristically human abilities, such as emotion, generalization of knowledge and its application elsewhere or the ability to predict and make plans based on past experience, the stronger (more general) it is. The narrower its scope, the weaker it appears in comparison.[15]

We have not yet reached *strong* AI, nor are we close, and not because of lack of effort. The development of such artificial systems is not an exclusively technical problem, but also neurological and philosophical. Philosophically speaking, the development of a truly strong AI carries great significance, specifically for the Philosophy of Mind (see Turing's Test & Searle's Chinese Room, the debate is still ongoing).

On the other hand, we do have what isn't strong AI. Defining these AI systems as *weak* implies that they are not able to perform extraordinary tasks, which couldn't be further from the truth. It would be more appropriate to define them as narrow AI since they are applicable to a specific setting or domain. Narrow AI can be found in self driving cars, social media algorithms and chatbots.

Machine learning is a form of narrow AI and a branch of computer science that focuses on the use and development of algorithms that use *data* to imitate the way humans learn, gradually improving their accuracy.[16] It was defined in the 1950s by AI pioneer Arthur Samuel as *the field of study that gives computers the ability to learn without explicitly being programmed.*[3] These algorithms are indeed able to learn and adapt, drawing inferences from data, without following explicit instructions, that is without being programmed to do so.[17]

To develop a successful ML system one must start from *data*, as the above definition suggests. A ML model is chosen and specific data gathered and prepared to be used as a training set (the set of information the model will be trained on). The larger the set, the better the preparation. The model trains itself on the supplied set to find commonalities and patterns among the data or to make predictions, for instance. Sometimes human intervention (e.g. parameters change) is beneficial to allow the model to reach more conclusive results.

Part of the data is reserved to be used as evaluation data in order to assess how the ML model performs to other information. Once the model is considered to be trained, it is exposed to the evaluation set. If the outcome is positive, then the model can be employed on real sets of data.

The **function** of successful ML systems can be manifold. It can be **descriptive** (the system uses the data to explain what happened), **predictive** (the system uses the data to predict what will happen), or **prescriptive** (the system will use the data to make suggestions about what action to take). [16]

There exist **three main types of ML models** (see Figure 2):

Supervised learning. These models are trained on labeled data[2], allowing the algorithms to learn and improve over time. The system is given an input and a specific output is requested. For instance, a model can be given a labeled picture of a banana ("Here, what I am showing you in these pictures are bananas") and asked to tell the banana apart from other things (a pipe for example). With time the algorithm can learn ways to accurately identify and differentiate bananas from other objects by itself.

Unsupervised learning. These models are trained on unlabeled data. An input is given but no specific output is requested. The algorithms may be able to find patterns, similarities and ways of classifying the data that were not explicitly asked for. The model might be shown many different pictures of fruits and learn to identify that some are "bananas", "apples" and so on, although it will not necessarily know how they are called.

Reinforcement learning. These models are usually employed in multi-step processes and work through trial and error, establishing a reward system. The required output is not disclosed to the system,

[2]Data that has been tagged with one or more labels identifying certain properties or characteristics, or classifications or contained objects

but when the model makes the right choice, it is given a reward, learning the right actions to take. Reinforcement training can be found in self-driving cars or virtual players (Chess, Go, etc.) among others[3] and is the closest model to the ultimate goal of strong AI.

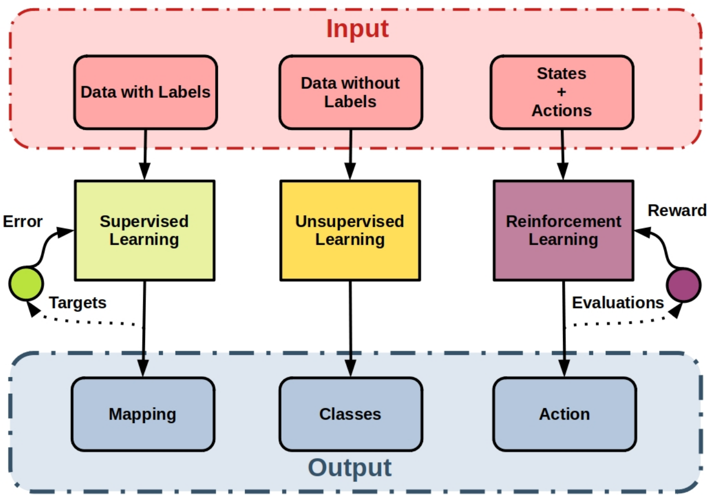

Figure 2: Types of ML models.[2]

It is important to understand that these different learning models have a specific domain over which they are efficient and on which their application makes sense. ML is not always useful. The map on Figure 3 developed by MIT researchers may aid in understanding this concept.

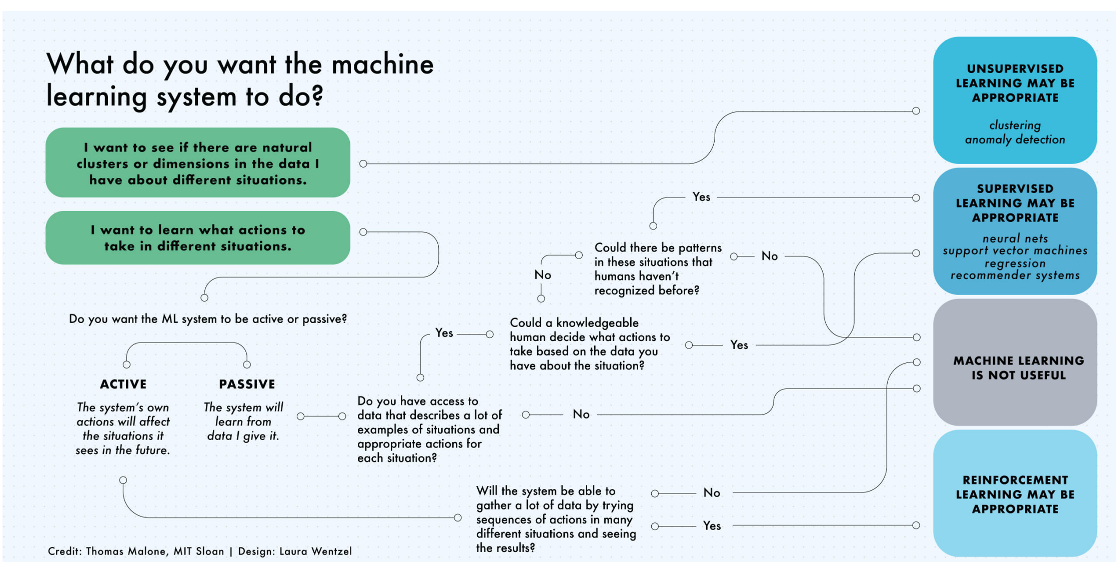

Figure 3: Applicability of ML.[3]

Aside from the classical structure embedded in these typologies (Figure 2), there are two additional ML **structures** that are of great importance, both because of their potential and their increasingly widespread application: **Deep Learning (DL)** and **Artificial Neural Networks (ANN)**.

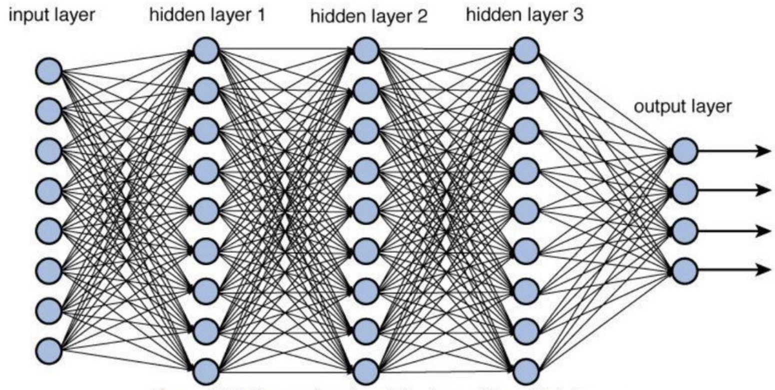

(a) Layers of DL system.[19]

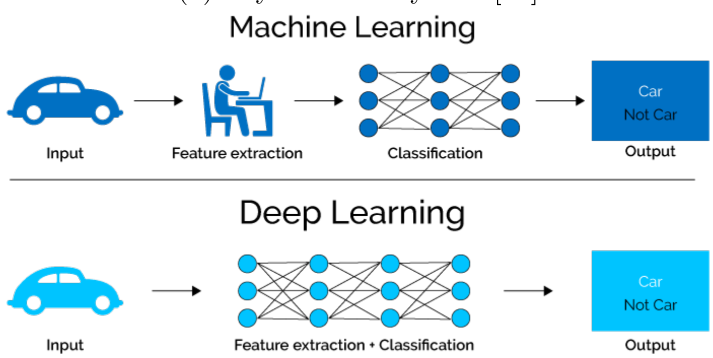

(b) Classical vs Deep ML structures.[20]

Figure 4

Deep learning owes its name to the many layers that conform the entire system. The larger the number of layers, the deeper the system. While in "classical" ML scenarios the system is composed of an input, a middle layer (where the magic happens) and an output, in DL systems there are many more layers between both extremes. The output of the first layer is fed as an input to the second layer and so on, until the final output is given (Figure 4a).

DL is most often used with supervised or semi-supervised algorithms. While "classical" ML structures attempt to extract information from a large set of pre-processed and structured data (with experts needing to formulate the rules the machine will learn on), DL automates the feature extraction part, eliminating some human intervention and enabling the use of larger data sets (Figure 4b). DL might be thought of a "scalable machine learning" as noted by Lex Fridman in one of his MIT lectures.[18]

On the other hand, an **artificial neural network** represents the structure of a human brain modeled on a computer. It consists of *neurons* organized into layers, frequently millions of them, connected into one system, making it extremely successful at analyzing and even memorizing information.

ANNs are comprised of node layers, containing an input layer, one or more hidden layers, and an output layer (see Figure 4a). Each node, acting as an artificial neuron, connects to another in the subsequent layer. If the output of an individual node is above a specified threshold value, the node is activated and it sends data to the next layer of the network. Otherwise, no data is passed along. A neural network that consists of more than three layers (inclusive of the inputs and the output ones) can be considered a **deep learning algorithm** or a **deep neural network**.[16]

2 Machine Learning in Finance

An important consideration to make is that ML is most often associated with statistical learning, although this need not always be the case. In finance, however, ML is fundamentally correlated to statistical modelling, to the point where some classical statistical methods (such as logical regression) are vastly employed in ML solutions and - although statistical methods by themselves are not ML - the boundary between the two of them is not always well defined (some financial statistical methods employed in the 90s are now considered ML).[7]

2.1 Looking back

The origins of applying advanced mathematical methods to the financial sector goes back to 1900, with the publication of Louis Bachelier's Theory of Speculation (*Théorie de la Spéculation*), one of the first works to explore the use of advanced mathematics to evaluate stocks. In fact, a considerable amount of the financial mathematics and statistical modelling that is still in use nowadays was first developed in the early 20th century and some argue it was the beginning of primitive AI in Finance.[21]

The initial development of AI systems in the 60s was not particularly concerned with finance, but during the 80s and 90s many AI financial solutions proliferated and were commercialized, giving rise to a new found interest in the technology to the point where two thirds of the Fortune 1000 companies had one or more AI project in development.[21][22] In particular, the 80s were "dominated"/ more concerned with the rise of Expert Systems [3], while in the 90s the attention shifted towards the possibility of using Fraud Detection algorithms.

However, in the 90s it became evident that most of these solutions failed to achieve the expected results and meet reality, perhaps because of the technological impediments in storage and processing capacity or because of the excess of power and flexibility they provided, which in Finance is most usually not a good thing. A helpful analogy provided by Dr Halperin is that of a novel driver who, used to driving a Toyota, suddenly finds themselves in a Ferrari. Chances are they will end up crashing because of the excess of power of the car. [7]

In any case, AI systems overpromised and underdelivered, leading to the so called *AI winter*, a period of reduced funding and interest in the subject (although not complete disappearance) that lasted almost a decade.

Enthusiasm was recovered in the 2000s, most probably in relation to technological advances in computer capacity and large and growing availability of data within financial services. By 2006 algorithmic trading[4] was present in a third of the European and US market.[23] High frequency trading[5] (HFT) accounted for ∼65% of the US equity trading volume in 2009.[24] In July of 2010 the first case of the use of ML in investing was reported by the Wall Street Journal.[25]

In the 2000s and early 2010s there were several ML breakthroughs related to the shift towards DL, the so-known as *Deep Learning Revolution*, although not everyone agrees on designating it as a real

[3] A computer system which emulates the decision-making ability of a human expert, designed to solve complex problems by reasoning through bodies of knowledge, represented mainly as if–then rules. Nowadays, they are obsolete.

[4] An order execution method based on automated pre-programmed trading instructions accounting for variables such as time, price, and market volume

[5] A type of algorithmic trading characterized by large number of transactions in fractions of a second, which analyzes the market and executes orders based on market conditions.

revolution. The main achievements consisted in getting old networks to work and outperform their old selves by exploiting much larger amounts of data, having arranged them in a Deep architecture. It was remarkable progress but mostly driven by technology rather than by fundamental science.[7]

2.2 Today

The adoption of ML in the sector has been increasingly *deepening* for the last decade, and its application today (whether successful or not) is present in almost all corners of the subject, although compared to other fields its presence is still marginal.[8] Figure 5 gives a few example on the use of automated learning.

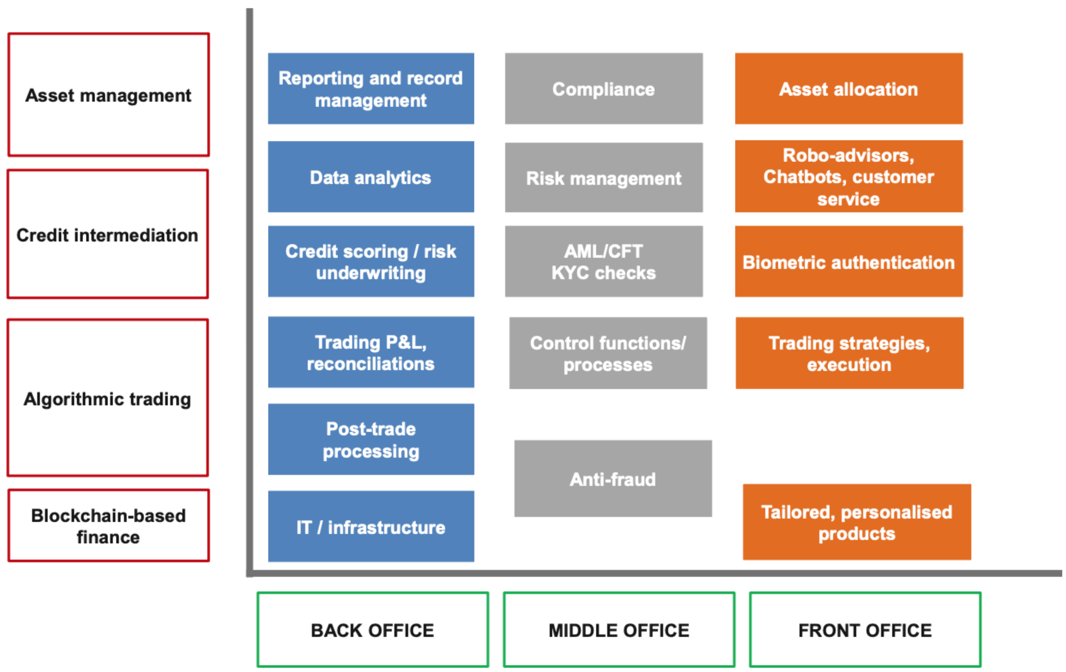

Figure 5: Some applications of ML in financial activities.[4]

The predominant sectors for ML employment in finance are **Fraud Detection**[6], **Credit Scoring**[7] and **Intra-agent Trading**, mainly because of the large abundance of data that can be found in their regard - though not without lack of criticism and challenges - as well as the partial automation of customer service (Chatbots, Robo-advisors, etc.). Arguably though, the most appealing application lies within **investments**, **financial markets** and **Asset Management**[8].

For the investment community, information has always been key and data has been the cornerstone of many investment strategies. Feeding ML models with data could provide asset managers with recommendations that influence decision-making around portfolio allocation or stock selection, depending on the type of ML technique used.[4] However, the simplicity of the concept is pretty much only apparent

[6]The concept lies in the assumption that fraudulent transactions have specific features that legitimate transactions do not. ML can detect patterns in financial operations and decide on its legitimacy much more effectively than humans by processing mountains of information faster and spotting patterns that might seem unrelated or go unnoticed by humans.(Intellias)

[7]There exist countless variables that might predict an applicant's ability to pay back their loan, and ML is good at finding patterns within large data sets, factoring in data points that are as of yet unknown to predict a borrower's likelihood of paying back their loan (Emerj)

[8]The management of investments on behalf of others.

since the employment of ML in finance is fundamentally different from other technological sectors for a few reasons.

Igor Halperin is an AI/Quant Research Associate at Fidelity Investments.

His research focuses on using methods of Reinforcement Learning, information theory, and physics for financial problems such as portfolio optimization or dynamic risk management. Until 2020, Dr Halperin worked also as a Research Professor of Financial Machine Learning at NYU Tandon School of Engineering. Before that, he was an Executive Director of Quantitative Research at JPMorgan, and a quantitative researcher at Bloomberg LP. Dr Halperin obtained his Ph.D. in theoretical high energy physics from Tel Aviv University, and his M.Sc. in nuclear physics from St. Petersburg State Technical University.

For him ML is inherently related to finance: in fact, the first time he heard of the term Machine Learning was shortly after he joined the financial sector, around 1999.

2.2.1 Peculiarities, Hazards and Protagonists

First and foremost, **financial data** is inherently different from other types of data found in physical or empirical situations. Contrary to popular belief, financial data is not *Big* but rather *Small*, excluding some particular exceptions,[7] making the training stage very challenging. This by itself limits tremendously the applicability of ML but the peculiarities go way deeper. Financial data is also characterised by its non-stationarity and its low signal-to-noise (S/N) ratio,[7] meaning its properties depend on the time at which the data is observed and that the data itself is very noisy (accompanied by interferences that obscure the signal). Moreover, a great deal of the data available nowadays is in an unstructured form (∼80%), preventing its immediate use for most ML solutions.

In a way, it is easier to build a self-driving car than it is to develop a money earning algorithm since in the physical world rules don't change, empirical laws stay the same. In financial markets, on top of the non-stationary and noisy data, because of and the phenomenon of arbitrage[9] chances are that any system that works will probably stop doing so at some point, since the market will eventually learn from it.[8]

It can be argued that an analogous data complexity can be found in the medical field, judging by the above specified parameters but also by the added difficulty (in comparison to other areas) in taking into account the effects of external factors within the data itself (think of politics in finance or environmental contamination in medicine), as Dr Halperin suggests. In both sectors, indeed, many believe ML has ended overpromising and underdelivering slightly.

The data scarcity has led many to think that the evolution of ML will be marked by the advent of ***synthetic data***. Synthetic data is non-real data, in the sense that it was not gathered from real life, but is generated ad hoc (through algorithms) with the same features that real data possesses. Outside of the financial environment, synthetic data production can be easily found in the website `https://thesecatsdonotexist.com/` for instance. As the name suggests, the cats shown in the images on the portal do not exist (see Figure 6), they are artificially generated by an AI system that has

[9] The simultaneous purchase and sale of the same asset in different markets in order to profit from small differences in the asset's listed price.[1]

abstracted the features that make a cat seem like one and has learned to replicate them. This new data can in turn be used as training set for the algorithm to improve.

Figure 6: Cats that do not exist (screenshot from website).[5]

In fact, without an appropriately large data set, the algorithms could reach incomplete or inappropriate conclusions that may lead to grave consequences were they to be employed in a real setting: imagine you develop an algorithm that is able to guess the number you have in mind. However, you only give the system numbers from 1 to 10. Of course, the system will be able to discern several apparent solutions (all of which will give the correct answer for the limited numeric domain), but the instant you increase the numbers from 10 to 500'000 most of those methods will probably fail since they were not real solutions .[8] In the absence of real data, synthetic one could assist.

Within the financial ecosystem, however, the generation of synthetic data is more complex as Dr Jaureguizar states. In finance, these generative models need to simulate the statistical distribution in which real data is presented and be able to capture the non-stationarity and the external interferences that are so characteristic of it.

The challenges in data instigate the need for the development of more robust[10] and less data hungry methods[7]. Physical limitations play a central role too, since most of these methods are incredibly powerful but at a very high energetic and monetary cost and take a non-trivial amount of time to be trained. Storage and processing capacity, although infinitely better than 10 years ago, are still not enough sometimes. The development of quantum computing might provide a solution[8] but in the meantime one must carefully think about the associated *cost vs profit* of the algorithmical implementation[7] (if a ML operation comes at a cost of e.g. \$50'000 but the profit made from it barely reaches \$1000, it was clearly not worth it).

A common criticism of financial ML implementation is the limited **interpretability** it provides.[7] Oftentimes ML methods tend to act as a *black box*, an *oracle* whose words must be followed blindly

[10]The capacity of performing without failure in a wide range of conditions.

without aspiration to understand the reasons behind its imperative. In finance (an other sectors such as medicine) this can lead to feeling of overconfidence and a false sense of security. - non-trivial amounts of money have been lost for these reasons. Blind obedience is not enough and the people employing these methods must have at least a basic understanding of why to follow the suggestions of the ML system. Moreover, ML presents an added debugging difficulty, since errors are not always evident and their appearance not always immediate.[8] In an empirical setting, you could confidently verify the validity of your solution if you were to design an experiment and test it. In finance however, considerable amounts of time can go by without noticing fundamental mistakes until it's far too late.

Explainability and the concept of **Explainable AI** (**XAI**)[11] are crucial for the correct application of ML to the financial sector (although some argue even XAI is not that easily comprehensible).

Carlos Jaureguizar is CEO of Roexia AI Consulting, a ML lab whose application is mainly centered around the financial sector.

The lab was born as a natural evolution a previous business centered around financial quantitative systems, Noesis. Dr Jaureguizar defines the progression from quantitative to AI businesses as the natural evolution of the market. Until 2019 Dr Jaureguizar was president of the Spanish Institute of Technical and Quantitative Analysts. He obtained his Ph.D. in applied economics from the Universidad Rey Juan Carlos and his M.Sc. in banking and financial markets by the Universidad Autónoma de Madrid.

Although he first heard of AI & ML around the early 2010s, it wasn't until 2016 that it became a daily reality for him. He now provides state of the art technology and is one of the first to develop AI embedded commercial solutions for the financial sector in Spain.

A second common and very relevant criticism is that ML, by definition, cannot predict what's unpredictable. Because of so it may appear useless when faced with drastic unforeseen circumstances (e.g. the Covid-19 pandemic). In finance the appearance of such unpredictable events whose consequences might be extreme is known as a ***Black Swan*** because of its rarity. The longer your forecast horizon (in the future), the more important these phenomena become.

While it is absolutely true that ML lacks the ability to predict Black Swans, what it has is the ability to react to the abrupt change of conditions caused by the advent of such an event itself and adapt its strategy in consequence - by liquidating positions for instance. A very insightful real life example was provided by Dr Jaureguizar during the interview conducted this past August:

"Emilio Villota, a spanish F1 driver, talks about the time when an accident occurred right after a corner during a race. Obviously he didn't know and had no way of knowing, but in spite of this, he stopped before the corner and was able to avoid getting involved in it. When he was asked how he knew an accident had happened he simply answered that he didn't but that every time he drove along the straight of the circuit he could see many white dots on the laterals (of course, these were people's faces that could not be clearly discerned at such high speeds), but the last time he drove by, the white dots had disappeared and been substituted by black ones, meaning that people were all looking at something. If everyone was looking at the same thing, he said, something must had happened, so he braked before the

[11]Subset of AI in which the results of the solution given by the system can be understood by humans

corner. And that is exactly what ML can do, step aside if needed."

In contraposition, were something amazing to happen, because of the abnormal change in conditions, the ML algorithm would still choose to step aside (since the focus of risk management lies on minimizing rarities) and lose any potential competitive advantage derived from it. A calculated risk.[8]

Sandra Nieto is Chief operating officer (COO) of Robexia AI Consulting.
She obtained her M.Sc. in Advertising and Marketing by the Universidad Complutense de Madrid on 2016.

Ms Nieto is a key piece in the implementation of Robexia business strategy. Her implication in Machine Learning comes as a natural evolution of the business model of Robexia, which entails her expertise towards the product transformation.

Her extensive experience within the field in which Robexia operates, is essential to design the marketing strategy of the developed trading products, providing customers with a qualitative approach to Machine Learning.

No one knows exactly how widespread ML truly is within the industry. What is certain is that there is an increasing demand for ML solutions that has not been met by a corresponding offer yet. Smaller labs and specialized startups are beginning to increase in numbers but the use of ML is most often reserved to larger asset managers or institutional investors who have the capacity and resources to invest in AI technologies, creating a potential blockage for the adoption of the methods by smaller agents.[4]

The phenomenon of restricted participation may continue until the access to ML tools becomes "universal". When this point is reached, however, the employment of the same models by a large number of institutions could lead to one-way markets and herding behaviors which could create severe liquidity and stability issues for the whole system.[4]

Theoretically, there are no areas to which ML could not be applied to, but it must be done with caution. It is imperative to understand that ML is a tool and must be used as such. Expecting ML to adopt the role of *theory creator* is a fundamental mistake that must be avoided at all costs; ML is a tool that can only be as good as you define it to be (it is your duty to define the task, define the method and put into use the software and hardware).[7]

Not all ML methods are useful in all situations and not all situations require the use of ML. M&A[12] (Mergers and Acquisitions) for instance is still completely humanly driven because of its high complexity.

Although understanding this might seem trivial to some, it is not unusual to see people promising the Moon through means of ML and for people that have no expertise in the subject it is not always easy to see through them. As explained by Ms Nieto, when commercializing ML products one usually focus on the solution they provide rather than on the technical specifications, often through examples, since explaining the guts of the algorithm unfortunately scares most off. That is why ML experts must fit a specific profile to ensure the solutions they provide are reliable and can be called back for maintenance if embedded in a larger product.

[12]Consolidation of companies or assets through diverse financial transactions, including mergers or acquisitions.

The Conway diagram (Figure 7), developed in the late 2000s, shows the superposition of skills that such an expert should carry. It can be argued that nowadays, with the several specifications within ML the correct profile would be the central (”Data Science”) and not the top (proper ”Machine Learning”) one, since the peculiarities of the sector in which ML methods are to be employed play an incredibly relevant role. In real life, however, these profile are not strictly defined and the relative novelty of the subject makes academic preparation challenging to find.[9] We have come a long way but much more is yet to come.

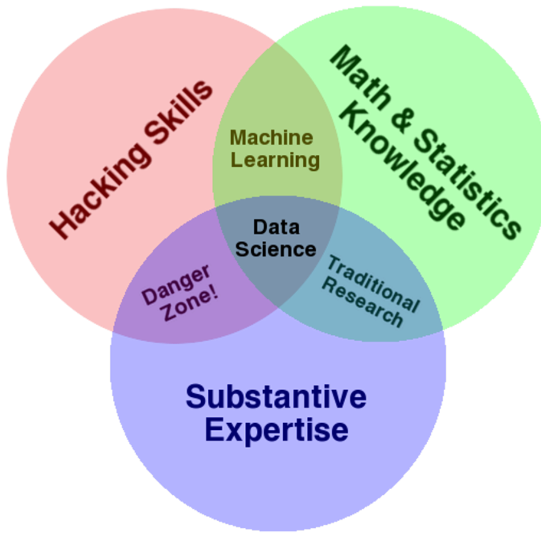

Figure 7: Skill profile of the ML expert.[6]

3 Conclusion

It seems reasonable to conclude, as Dr Halperin claims, that Machine Learning has not yet reached the point where it plays a pivotal role within the financial sector, and it may be argued that it is not necessarily destined to do so.

Undoubtedly, Automated Learning has been an extremely useful complementary tool in many areas but by itself it lacks the capacity of flipping the picture entirely. ML can only be a part of the solution.

The particularity of the tool, both in technicality and applicability, calls for a cautious and supervised employment of ML. A thoughtful study should accompany its adoption to ensure its correct exercise. New regulation will need to come into order to ensure accountability and avoid extreme competitive disadvantages that could alter the system dangerously to occur.

As it usually happens in tech, the progression of ML systems will most probably be exponential although staggered, as Dr Jaureguizar suggests. As of right now we have not yet reached the steepest point of its growth but many think, with the the arrival of new talent and improved technology, we will be able to see it in the coming years. In fact, rather than the *Revolution* many believe to be currently taking place, Dr Halperin suggests we have been witnesses to a fruitful and steady *Evolution* that will continue to expand its influence in the following years. Dr Jaureguizar hints, however, we might be but a few steps away of the inflection point in which the proper *Revolution* could begin.

References

[1] Investopedia, "What is finance?." URL https://www.investopedia.com/ask/answers/what-is-finance/.

[2] "Supervised vs unsupervised vs reinforcement learning."

[3] S. Brown and M. S. M. School, "Machine learning, explained," April 21 2021. URL https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained.

[4] "Artificial intelligence, machine learning and big data in finance opportunities, challenges and implications for policy makers."

[5] "These cats do not exist." URL https://thesecatsdonotexist.com/.

[6] D. Conway, "Conway diagram," 2010.

[7] I. Halperin, 2021. Interview by Carlotta Berto Salmeron.

[8] C. Jaureguizar, 2021. Interview by Carlotta Berto Salmeron.

[9] S. Nieto, 2021. Interview by Carlotta Berto Salmeron.

[10] W. E. Forum, "Fourth industrial revolution." URL https://www.weforum.org/focus/fourth-industrial-revolution.

[11] U. of South Alabama, "What is economics, what is finance, and what is real estate?."

[12] J. McCarthy, "What is artificial intelligence?," 12 November 2007.

[13] "Dartmouth conference," 1965.

[14] J. Howe, "Ai in england," 1994.

[15] K. Walch and Forbes, "Rethinking weak vs. strong ai," October 4 2019. URL https://www.forbes.com/sites/cognitiveworld/2019/10/04/rethinking-weak-vs-strong-ai/.

[16] I. C. Education, "Machine learning," July 15 2020. URL https://www.ibm.com/cloud/learn/machine-learning.

[17] O. O. E. Dictionary, "Definition of machine learning,"

[18] L. Fridman, "Mit lecture," January 11 2019. https://www.youtube.com/watch?v=O5xeyoRL95U.

[19] "Deep leaning hidden layers." URL https://towardsdatascience.com/training-deep-neural-networks-9fdb1964b964.

[20] "Classical vs deep architectures." URL https://semiengineering.com/deep-learning-spreads/.

[21] S. Swaine-Simon and Medium, "The history of ai in finance," URL https://medium.com/district3/the-history-of-ai-in-finance-7a03fcb4a498#bcb7.

[22] J. Durkin, "Expert systems: History and applications," 2002.

[23] "Aite group." www.aitegroup.com.

[24] R. Iati, "The real story of trading software espionage," 2009.

[25] S. Patterson, "Letting the machines decide." https://www.wsj.com/articles/SB10001424052748703834604575365310813948080.