DEEP LEARNING IN FINANCE: FROM IMPLEMENTATION TO REGULATION

Louis Bertucci
Marie Brière
Olivier Fliche
Joseph Mikael
Lukasz Szpruch

EXECUTIVE SUMMARY

Deep Learning in Finance: From Implementation to Regulation

• Despite important theoretical questions that remain to be solved, Artificial Intelligence and Deep Learning are being increasingly used in the Finance and Insurance sector.
• Beyond straightforward data analytics, decision models are being implemented with Deep Learning. These algorithms cannot be used blindly. The understanding of the underlying problem is key. Humans, engineers or mathematicians, are essential.
• One trendy application is the use of Deep Learning (specifically GANs) to generate datasets. In finance, data are often scarce and having the possibility to generate new data (similar to an original dataset) can be decisive.
• In many applications, explainability of Artificial Intelligence is critical to protect consumers.
• Explainability is not a one-size-fits-all concept, and several degrees of explainability may have to be reached. Explainability to non-specialists is an additional challenge.
• Biais in the learning data is critical to assess because biases will be reproduced by the algorithm, and lead to unexplained discriminations.
• The role of regulatory agencies will be crucial to protect consumers while allowing innovation. There is currently no unified regulatory framework. The European Commission's Artificial Intelligence Act (draft proposal in April 2021) lists prohibited artificial intelligence practices and defines high-risk application areas for which they identify requirements (risk management system, data governance, technical documentation and record keeping, transparency, human oversight, accuracy, robustness and cybersecurity).

Louis Bertucci, Institut Louis Bachelier, France
Marie Brière, Amundi Asset Management, France, Université Paris Dauphine-PSL, France and Université Libre de Bruxelles (ULB), SBS-EM, CEB, Belgium
Olivier Fliche, ACPR, France
Joseph Mikael, EDF R&D, France
Lukasz Szpruch, University of Edinburgh, School of Mathematics, Edinburgh, UK and The Alan Turing Institute, London, UK

INTRODUCTION

From the hedging of complex portfolios to the selection of features and models, finance often deals with high dimensional problems which are hard to tackle with traditional tools. Artificial Intelligence, and more specifically Deep Learning, offer a new paradigm to address the numerical solution of such problems. Traditional finance problems such as trading under transaction costs or with liquidity constraints are well suited to the use of Machine Learning and Deep Learning. For retail finance, these can, for instance, simplify and accelerate consumers' credit access procedures, offer better and more personalized access to financial advice via robo-advisors, and allow more efficient interactions through the use of chat-bots.

As in many other industries, Deep Learning has proven very useful for approximating complex non-linear systems, but it also gives rise to new challenges that need to be addressed both in the implementation and regulation of such algorithms.

Over the past decade, AI research has kept extending the boundaries of what it is possible to achieve. Researchers around the globe are constantly finding new use cases and applications. Much of the research carried out on this topic in both academia and industry is possible through the democratization of machine learning libraries like TensorFlow and PyTorch. The cost of building and running a Deep Learning algorithm keeps falling over time and for this reason the use of AI in finance, together with overcoming the associated regulatory challenges, are likely to remain crucial.

1. IMPLEMENTATION CHALLENGES AND PERSPECTIVES

Deep Learning theory

Deep Learning algorithms have great potential to help finance professionals understand market dynamics and game theoretical problems encountered in various situations. But one of the major challenges of Deep Learning is the lack of theoretical understanding and lack of results surrounding the training and convergence of the most used algorithms. Theoretical results essentially relate to the existence of an architecture that is able to approximate a given function, whereas the real challenge is the learning process[1]. Many fundamental questions remain beyond what is known, eg. when will the learning process end and when will Neural Networks finally converge to the solution and crucially to which solution out of many possible ones it will converge? By way of illustration, the goal of the learning process is to minimize a loss function, but given the large number of parameters, several local minima may exist , and it usually cannot be proved that the algorithm has converged to the global minimum[2]. The lack of understanding of the learning process also extends to the speed of convergence. It turns out that most of the time (when the hyper-parameters have been well defined), the learning process converges relatively quickly to a solution (sometimes several orders of magnitude faster than traditional numerical tools). Why this happens and under what conditions is still a puzzle researchers working on theoretical foundation of machine learning are wrestling with.

The lack of theoretical results is a fundamental concern, because it can erode trust in the use of Deep Learning algorithms. Given the critical nature of those algorithms' use cases in the finance industry, we need to be highly confident that the solution obtained after the learning procedure is the one desired.

Operational use of Deep Learning algorithms

Despite these theoretical issues, Deep Learning algorithms are increasingly being used to solve concrete practical problems[3]. And in many applications, the training process works remarkably well. In the financial industry, deep learning is used for optimal execution, trading strategies, the hedging of complex risks, fraud detection, portfolio advising (robo-advising), and textual analysis such as financial statement analysis and market sentiment measures. In the insurance sector, Deep Learning has been used, among other things, for risk management, client/risk classifications, underwriting of new clients, and fraud detection. The energy sector is also well advanced in terms of use cases. With the financialization of energy markets such as electricity, actors in the industry have started making extensive use of Neural Networks in their operations.[4]

[1] See for instance Cybenko (1989) "Approximation by superpositions of a sigmoidal function", Mathematics of Control, Signals and Systems 2 (4), Hornik (1993) "Some new results on neural network approximation", Neural Networks 6 (8), and Gelenbe, Mao and Li (1999) "Function approximation with spiked random networks", IEEE Transactions on Neural Networks 10 (1).
[2] There are in fact some results on the convergence of Neural Networks, but they often do not generalise well or they focus only on a specific architecture. See for instance Li and Yuan (2017) "Convergence analysis of two-layer neural networks with ReLu activation", arXiv preprint, or Jentzen and Riekert (2021) "Convergence analysis for gradient flows in the training of artificial neural networks with ReLu activation", arXiv preprint.
[3] See, for instance, Buehler, Gonon, Teichmann and Wood (2019) "Deep Hedging", Quantitative Finance 19 (8) or Charpentier, Elie and Remlinger (2021) "Reinforcement learning in economics and finance" Computational Economics 1-38.
[4] The research lab of EDF (the French traditionl electricity provider) has been quite active in this field. See, for instance, Fécamp, Mikael and Warin (2020) "Deep learning for discretetime hedging in incomplete markets", Computational Finance 25 (2).

In the electricity market, traders face very strong liquidity constraints, multiple sources of uncertainty (climate, intermittent production, etc.) and have to manage relatively large transaction costs, making it impractical to solve grid optimization with traditional optimization tools. Consequently the usual way to

handle such constraints is to considerably simplify the underlying problem. Deep Learning allows both liquidity constraints and transaction costs to be handled by means of a few lines of code in a well specified problem without simplifying the optimization problem.

Increasing the dimensionality of the problem without studying the underlying theoretical model is sometimes sufficient, but this naive approach also has certain drawbacks. Relying too much on Neural Networks can give rise to trust issues for the user of the algorithm, because in some instances there is no way to check whether the solution is correct. On the other hand, the naive approach allows previously complex problems to be solved in practice. Therefore operational teams have a strong incentive to use Neural Networks right away.

A better approach would be to use detailed knowledge of the problem to solve, and to study, the theoretical model that would allow a Neural Network to be used just for those parts of the numerical scheme that cannot be handled easily[5]. Indeed, Neural Networks are not models and they shouldn't be used as such. They are, however, extremely efficient high dimensional function approximators and interpolators. Such mixed methods highlight the need for human intelligence and knowledge about a particular problem as well as the need to be able to derive a numerical scheme that can leverage the power of both Neural Networks and mathematical models.

Human-machine interaction

It can be tempting to blindly use Neural Networks to solve any problem (which sometimes does work), but doing so would eventually lead to overexposure to the risk of misspecification and approximation errors. Human-machine interactions are key in the use of machine learning algorithms in several respects. First, a naive Neural Network is less likely to work than a trained numerical scheme designed by field specialists that relies on Neural Networks for only parts of the algorithm. Human design is thus really important.

But once a Deep Learning model has been designed and trained, it can provide new insights for the developers of this very model and specialists in the field. This approach has been extensively used in fields such as board games, with Deep-Mind being very successful in this regard[6]. For board games in general (Chess, Go, etc.), the point has been reached where Deep Learning models perform far better than human players. Indeed human players are learning from these algorithmic players in order to acquire new insights and strategies. But the situation is different in finance and economics due to the nature of the field. Chess is a two-player game and is the same whether played in Paris, France in 2022 or in 1956 in New York, USA. In finance, on the other hand, markets behave very differently over time and aggregates such as prices are the result of the interactions of thousands of individuals, whose behavior can change and adapt.

[5] See, for instance, Achdou, Lasry and Lions (2022) "Computing Mean Field Games in infinite dimension I : The example of the Krusell Smith model", forthcoming, and Achdou, Bertucci C. and Bertucci L., Lasry and Lions (2022) "Bitcoin Mining with Several Generations of Miners", forthcoming.

[6] See, for instance, DeepMind's AlphaZero or AlphaGo for examples of such algorithms.

Finally, the question of human responsibility is crucial. Used well, a Neural Network can help a risk manager handle complex portfolio risks by generating metrics and measures, while the ultimate responsibility lies with the human decision-maker. But the issue of responsibility is essential for regulators, as some actors may try to pass on their responsibility to algorithms. In this context, the explainability of Artificial Intelligence is particularly important, especially when models are used without human involvement (in High Frequency Trading or automated client insurance underwriting). Explainability is also important for reinforcing users' trust in this new tool. This idea will be developed in more detail in Section 2.

Data Generation Techniques

In contrast to the traditional use of AI in fields such as image or speech recognition, financial applications usually suffer from a lack of sufficient data to train the models. Indeed, in the case of stock returns for instance, there is an inherent limit to the number of trading days per year for a given company or index. This makes the training of some AI models such as deep hedging or algorithmic trading tricky and prone to a number of biases. Moreover, a large amount of data may be needed to perform stress testing and risk management. Finally, in banks and other financial institutions, data sharing between different business lines may be limited. Datasets are often stored in silos due to regulatory requirements. It may therefore be useful to create synthetic datasets that share the same properties as real data, while respecting privacy requirements.

To overcome such data scarcity, AI can be used to generate more data[7]. This may seem counterintuitive because data is needed in the first place to train the AI algorithm. However, Generative Adversarial Networks (GANs) can be used to generate a new data sample that is intended to have properties similar to the original time series. More precisely, a GAN will produce a data sample that will be hard to differentiate from the original data based on certain statistical properties.

As usual when using AI, there are risks associated with the data generation process. Although ideally one would design a model-free market generator, certain assumptions still need to be made and the design required to make GAN works is a challenging task. In particular, even if the statistical properties of the original data are matched, the GAN will reproduce all its biases. While standard in all AI applications, this is especially relevant for data generation because this data will be used by another algorithm to make a decision or generate insights.

Human-machine interactions are also essential in this process. When applied to risk management, risk simulations do not have to rely solely on these new types of market generators, and traditional stress tests should still be performed in association with the scenarios generated by a GAN[8]. Human oversight is important, because pure AI algorithms suffer from too many biases and too much uncertainty, especially when applied to critical applications like simulations of banking risks and the pricing of derivatives exposure.

[7] See for instance Buehler, Horvath, Lyons, Arribas and Wood (2021) "Generating Financial Markets With Signatures", available at SSRN 3657366 or Ni, Szpruch, Sabate-Vidales, Xiao, Wiese and Liao (2021) "Sig-Wasserstein GANs for Time Series Generation", arXiv preprint arXiv:2111.01207 ; Remlinger, Mikael, Elie, Conditional Loss and Deep Euler Scheme for Time Series Generation, proceedings of AAAI 2022.
[8] For an application of GAN to risk simulation of trading strategies see, for instance, Lezmi, Roche, Roncalli and Xu (2020) "Improving the Robustness of Trading Strategy Backtesting with Boltzmann Machines and Generative Adversarial Networks", available at SSRN.

2. HOW SHOULD ALGORITHMS BE SUPERVISED?

The use of AI algorithms is growing rapidly in the financial sector. As we have seen, there are a number of interesting applications and use cases for machine learning algorithms but using these algorithms can also increase certain risks. Deep learning algorithms, which rely on an enormous number of parameters, can behave unexpectedly. For instance, they can be very sensitive to small changes in inputs[9] and arbitrarily change their predictions. More generally, AI algorithms may exhibit unintended bias or discrimination regarding certain groups. Thus, although they perform well over a wide range of tasks, their underlying mechanisms are very different from those used by human beings and are not always fully understood. Hence their usage may result in unanticipated financial and operational risks for the financial institutions concerned.

The importance of explainable algorithms

Given this situation, the explainability of algorithms is essential, both for users, for whom explanations make the system easier to operate and promote trust, and for regulators, who will have to supervise their use in order to prevent excessive risk exposure[10].

Two aspects of machine learning models may need to be explained: (i) the learning process, where one needs to understand the algorithm, the objective function, the training and test data, the model's tuning parameters, and (ii) the inference process, where the trained model makes predictions using real data. In the second phase, the question arises as to what are the main factors leading to the decision, and how the decision would be changed if one of these factors were altered. Finally, it may be useful to visualise the decision process synthetically.

Explainability can be investigated prior to the algorithm's learning phase (through an exploratory analysis of the data allowing, for example, dependence on sensitive variables to be measured), during the learning process (through joint prediction/explanation models), and subsequently in the post-modelling phase. A distinction can also be made between methods seeking global explainability (i.e. the ability to explain the functioning of an algorithm as a whole), and those providing local explainability, which make it possible to explain a particular algorithmic decision by quantifying locally the influence of the predictive variables on the decision (using, for example, the LIME, SHAP or SLIM[11] methods or counterfactual explanation methods[12]).

All these methods have their relevance, since explainability is not a one-size-fits-all concept. As a matter of fact, several levels of explanations (including "interpretations") can be distinguished depending on the person for whom the explanation is intended. In particular, providing a simplified but accurate explanation of an algorithm's result to a non-specialist is an additional challenge, which should not be underestimated. In a regulated sector such as the financial sector, the need to provide explanations will also depend on the risks associated with the process and/or regulatory constraints.

[9] This type of modified input is called an «adversarial example».
[10] Dupont et al. (2020), "Governance of Artificial Intelligence in Finance", ACPR discussion paper.
[11] For example, the SHAP method makes it possible to explain the influence of each predictor variable on the predicted values, in a model-agnostic way.
[12] A counterfactual explanation answers the question: what is the smallest possible change in the predictor variables that would have led to a change in the outcome (e.g. credit acceptance)?

When choosing the algorithm, the designer will therefore have to verify that the explainability of the algorithm is compatible with the highest level of explanation needed.

Of course, it should be borne in mind that requiring an algorithm to be explainable and/or fair can affect its accuracy[13]. For example, such constraints can reduce the flexibility of the algorithm. Explainability can also compromise the security principle (some fraud and money laundering detection algorithms need to remain secret in order to prevent criminals from engaging in reverse engineering, thereby avoiding detection). Explainability have costs as well as benefits, which may depend on the application domain, its impact on society and the corresponding regulatory requirement. Beaudoin et al (2020)[14] define several categories of costs, in particular those related to the storage of data in dedicated registers[15] allowing the reconstruction of the decision afterwards, and the design costs of the explanation function.

All these considerations will have to be taken into account in the global costs-benefits analysis of an AI algorithm applied to a given use case.

Dealing with the risk of discrimination

Another issue related to the use of algorithms is their fairness and ethics.[16]

The Villani report (2018) highlights the need to define tools for assessing discrimination. Even though European regulations prohibit taking into account sensitive personal data (religious, political, sexual orientation, ethnic origin, etc.) without explicit consent or substantial public benefit, the fact that the 'sensitive' variable is not taken into account or is even removed from the learning data is not enough to ensure that the decision is unbiased. Sensitive information can be correlated with non-sensitive information (consumption habits, etc.) and thus contribute to the decision bias.
Learning reflects the training data. If the data is biased (for example, if it is not representative of the population, or if there is a structural bias in the population under examination), the algorithm may reproduce or even reinforce this bias, and thus promote discrimination. To avoid discrimination bias, statistical indicators need to be put in place to measure it and then remedy it either by changing the sample of data used in the algorithmic learning process, or by modifying the decision rule in order to favor the absence of a link between the decision and the sensitive variable. The challenge ahead is to reach a broad consensus on the adequate statistical indicators reflecting what fairness means for society and law– in other words to bring statistical concepts and social or legal values together.

What type of supervision?

There is currently no unified framework for the regulation of AI algorithms, but most countries have issued general recommendations and principles (see, for example, the OECD principles (2019) or the European Commission white paper (2020), which are broadly based on the following five principles: reliability, accountability, transparency, fairness and ethics[17].

[13] Even if the explanation is constructed independently of the algorithm, the requirement for explanation may lead to the rejection of a solution that is difficult to explain.
[14] Beaudoin V. et al. (2020), 'Flexible and context-specific AI explainability: a multidisciplinary approach', Technical report, ArXiv.
[15] This requirement may also potentially conflict with legislation that limits the storage of personal data (RGPD, 2016).
[16] Besse P. et al. (2018), "L'IA du quotidien peut-elle être éthique?", Statistique et Société, 6(3).
[17] For more details on the regulatory debates and the mathematical tools that could be developed, see the replay of the conference: https://www.institutlouisbachelier.org/en/multimedia-2/ilb-replay-2/webinar-fair-deep-learning-in-finance-from-implementation-to-regulation-replay/.

The EU's General Data Protection Regulation (2016) already gives users of algorithms the right to an explanation of the outcome, as well as the right to challenge the decision and to obtain intervention from the individual responsible for processing.

However, the European Commission's Artificial Intelligence Act will form the basis of the regulatory framework of AI in the EU. The draft proposal (April 2021) lists prohibited artificial intelligence practices, defines high-risk application areas (health, finance, public services, transport) for which the principles listed above are translated into requirements (risk management system, data governance, technical documentation and record keeping, transparency, human oversight, accuracy, robustness and cybersecurity).

In the financial sector, this cross-sectoral regulation inspired by “product regulations” will have to be articulated with the sector's specific regulations and supervisory practices. Indeed, financial sector authorities will have a role to play as «market supervisory authorities» for «high risk» algorithms developed by financial sector institutions. In addition, European Supervisory Authorities are to publish guidance on AI in the financial sector by 2024 according to the EU digital finance strategy.

FaIR

INSTITUT
Louis Bachelier