THE CAPCO INSTITUTE

JOURNAL

OF FINANCIAL TRANSFORMATION

TECHNOLOGY

Applying artificial intelligence in finance and asset management: A discussion of status quo and the way forward

JUERGEN RAHMEL

WEALTH & ASSET MANAGEMENT

#51 APRIL 2020

THE CAPCO INSTITUTE

JOURNAL OF FINANCIAL TRANSFORMATION

RECIPIENT OF THE APEX AWARD FOR PUBLICATION EXCELLENCE

CONTENTS

INVESTMENTS

TECHNOLOGY

ESG

DEAR READER,

Welcome to edition 51 of the Capco Institute Journal of Financial Transformation.

The global wealth and asset management industry faces clear challenges, and a growing call for innovation and transformation. Increased competition, generational shifts in client demographics, and growing geopolitical uncertainty, mean that the sector needs to focus on the new technologies and practices that will position for success, at speed.

There is no doubt that technology will be at the forefront of a responsive and effective wealth and asset management sector in 2020 and beyond. The shift to digitization, in particular, will see the speeding up of regulatory protocols, customer knowledge building, and the onboarding process, all of which will vastly improve the client experience.

This edition of the Journal will focus closely on such digital disruption and evolving technological innovation. You will also find papers that examine human capital practices and new ways of working, regulatory trends, and what sustainability and responsible investment can look like via environmental, social and corporate governance.

As ever, I hope you find the latest edition of the Capco Journal to be engaging and informative. We have contributions from a range of world-class experts across industry and academia, including renowned Nobel Laureate, Robert C. Merton. We continue to strive to include the very best expertise, independent thinking and strategic insight for a future-focused financial services sector.

Thank you to all our contributors and thank you for reading.

Lance Levy, **Capco CEO**

APPLYING ARTIFICIAL INTELLIGENCE IN FINANCE AND ASSET MANAGEMENT: A DISCUSSION OF THE STATUS QUO AND THE WAY FORWARD

JUERGEN RAHMEL | Chief Digital Officer, HSBC Germany

ABSTRACT

Artificial intelligence (AI) and machine learning (ML) are gaining more and more traction in finance and asset management. But AI/ML is a complex tool that requires specific skills to be created, trained, and interpreted well for a given task. In this paper, we discuss some of the context parameters to be considered in order to apply AI beneficially in financial settings. We explore a matrix of use-cases, following the lifecycle of asset management and structured by the type of underlying AI technology. As AI requires human setup and interpretation, we briefly review the role of us "humans-in-the-loop" of AI implementations. Finally, the emerging field of asset tokenization promises to disrupt the conventional markets and market practices, opening up for a new field of AI applications to tackle the new way of trading and servicing securities. The AI game is on in asset management. Not to play is not an option.

1. INTRODUCTION

To learn about artificial intelligence, let's start with making a cup of tea

In 1825, the English scientist Michael Faraday began the annual tradition of delivering the annual Christmas lectures at the Royal Institution in London. The lectures present scientific subjects to a general audience, in an informative yet entertaining manner.

In 2019, the lectures focused on the topic of statistics, probability, and artificial intelligence (AI).[1] One of the live experiments was for the audience to give instructions to a pretend-robot making a cup of tea.[2] For those who are familiar with the English tradition of tea drinking, this process could be understood by a human through sentences as simple as "make us a cupper" or "put the kettle on". However, challenges arise if these instructions were targeted at a machine.

The live recording of this lecture shows that when the audience did not specify the size of the tea bags, the type of mug, the volume of boiled water that needs to be added, when to start and stop pouring the water, and how to add the milk, we are not able to achieve our objective of making a cup of tea.

As there are more and more "AI packages" around that can be simply downloaded and installed, we see a proliferation of AI use-cases and applications, often with limited success. Those unsuccessful scenarios are due to a lack of understanding that AI is actually not a simple tool. It needs knowledge and skill (and sometimes a bit of art) to choose the right type of algorithmic approach and to provide a proper data and learning environment for AI algorithms to be useful. It needs skill to properly apply AI to a use-case, as well as to interpret the results. Some of the main considerations for AI applications are described in this paper.

[1] https://bit.ly/38galPw
[2] https://bbc.in/2la0yht

1.1 Context, framing, and data flow

The above example shows the dependency of execution on precise instructions and also demonstrates one of the basic concepts of applying AI: the importance of framing in knowledge representation. This concept was introduced by Marvin Minsky, one of the founding fathers of artificial intelligence [Minsky (1974)]. We need to be able to provide the preconditions, the meaning, and usefulness of actionable options as well as the resulting post-conditions for a stepwise process like the one described. Without this piece of context, which is equivalent to "common sense" in the human world, the AI algorithms do not have sufficient information and fail to achieve the objective as intended. We humans have perfected the usage of common sense and common knowledge, as well as the degree to which we are assuming such knowledge to be present in other humans. This allows us to convey concepts with a minimal amount of factual information.

Another example for this was given by my AI Professor Michael M. Richter. In a working group meeting, he announced there was some news to share and he would demonstrate the framing and context idea by giving us two versions of information, both 100 percent factually correct:

First, he said: "Our dear colleague Frank will soon get married".

Then he said: "Our dear colleague Frank will soon get married for the first time".

Both statements are correct, however, why would he add the words "for the first time"? By adding this piece of information that, albeit true, is actually not necessary, our Professor triggered a totally different view on the communicated fact. In the first case, it is positive news, everyone was happy for Frank and wished to congratulate him. After the second statement, everyone was thinking about failed marriages, divorces, people separating, and marrying multiple times. The urge to congratulate Frank was rather dampened.

The conclusion to this can be that the high-quality capability of our intelligence is actually not based on processing as much information as possible, thus not being purely data-driven. The high performance of our intelligence might be the result of the creation of common concepts, abstraction, and generalization as well as communication and evaluation of only the really necessary additional parts of information. If this is true, then computational processing of sparse data together with a clever way of building of common concepts will be key to high performing "intelligent" algorithms.

1.2 Optimization and alignment of incentives

The other basic idea of AI is that its algorithms aim to optimize an objective, a target, a purpose. For this optimization task, we need to specifically formulate an algorithmic valuation function that aligns with the real-world-task of an AI application. Imprecise definition of this function will hamper the results. Examples for this can be taken from early attempts to create winning strategies in AI algorithms for automated playing of video games. In one attempt, the AI was tasked to play the game Tetris and given the capability to use all buttons that a human player can press. The objective coded into the algorithm was to maximize playing time, as a proxy for achieving highest scores. At first sight, we can assume that this is a valid incentive, as typically the longer the playing time, the higher the score. However, after stacking bricks in no orderly fashion as fast as possible (achieving a short spike in the score), the AI learned to hit the pause-button right before it would lose the game, thus "maximizing" playing time in a way that was not intended and that is not useful for the objective of achieving the high scores [Murphy (2013)]. As Murphy summarizes his result: "the only winning move is not to play" (in which movie did we hear this before?).

Applying this logic to AI in finance and asset management, we would subject ourselves to the same pitfall if we do not give precise instructions to, e.g., "optimize portfolio investment return within a given risk budget" or 'maintaining a certain portfolio segmentation' to fulfill the promises given to customers and regulations around the investment. We must start with a framing exercise, specifying the data and methods needed to achieve a specific objective, and let the machine carry out the number crunching functions as per our instructions and definitions. In other words, we must provide the objective together with the real-world framing and context, as it arises from the asset management task we are asking the AI to help solve.

1.3 Transparency and explainability of results

Many AI approaches, especially the very data intense ones, suffer from the so called "black-box problem" [see, e.g., Vontobel (2018) for a broader discussion of this topic]. This means that the complexity of the computational models, together with the vast amount of data going into these computations, make it outright impossible to understand or explain the reasoning the "black box" is going through to reach its solution. This can be criticized as potentially having hidden issues like unfairness, biases, etc., hence, there is a recent call to action to make AI more self-explanatory. However, it is not even clear what this could mean, since [see also Google (2019)]:

- Different audiences will expect different levels of explanation. For example, the customer inquiring a loan decision in terms of their personal data versus the mathematician expecting a sound mathematical explanation of the decision equations inside the "black-box".
- Explanations could be requested in real time by the actors or decision subjects, as opposed to situations where an auditor or supervisor would like to access the reasoning at a later stage.
- Different use-cases might warrant different qualities of explanation. For example, in an operational context of handwriting recognition, all that matters is the resulting accuracy of the algorithm, whereas in customer facing or even medical, life-critical settings it would be careless not to attempt to investigate the path of reasoning inside the decision algorithms and to be able to correlate similar cases to ensure quality, fairness, and non-discrimination.
- It is still a question of technical feasibility of generating meaningful explanations, especially in the large-scale neural network models of deep learning, where millions of data points are intertwined by millions of statistical training operations. There is simply no straightforward way to explain the resulting system (and if there was, it would probably not be necessary to take all the efforts of neural training anyway).

With these different perspectives on challenges for AI applications in mind, section 2 will provide an overview of existing and emerging application use-cases for AI in finance and asset management, section 3 will look into the role humans play in the context of AI applications, especially in finance, after the financial crisis, and in section 4 we present one particular, newly emerging use-case in the context of tokenization of assets – an example of the increasingly emerging cross-overs of new technologies.

2. AI IN FINANCE AND INVESTMENT MANAGEMENT

Artificial intelligence is entering all stages of the investment lifecycle

AI is a technology that promises a number of advantages, such as being capable of looking into vast amounts of data, assisting humans with decision-making, executing simple operational tasks all by itself, and being increasingly explainable, hence,

Table 1: Use-cases for AI in asset management

	CLIENT VIEW	MARKET VIEW	PORTFOLIO VIEW	MIDDLE/BACK-OFFICE OPERATIONS	RISK MANAGEMENT
NLP	Observe sentiment on social media, detect outreach opportunities	Compile Information from earnings reports, assess sentiment in Market comments. Support higher level instrument analysis and facilitate decision making		Observe conduct risk and morale of employees	Risk assessments based on analysis of texts and other unstructured information
BEHAVIORAL TREND ANALYSIS	Analyze own website traffic for insights on user behavior. Identify growth options, e.g., through nudges to induce client activity	Detect trends in industries, trade timing, technology, etc., that are induced by market participants	Extract "herd-movements" in market data for own portfolio as early chance/ warning indicators		Risk mining, i.e., crowd-sourced continuous internal risk assessments utilizing the swarm intelligence of the staff of a company to detect and identify risks. Those risks are anticipated by insiders of departments, industries, cultures, projects, etc.
AI-AND ML-BASED PATTERN RECOGNITION	Detect trading patterns of clients and client groups (by size, industry, over time, etc.)	Analyze unstructured data from alternative sources, like IoT-data (from e.g., vessels, goods in trade contexts)	Investment decision support based on large scale data views that humans can not conquer	Monitor for suspicious transactions, and trigger required actions	
	Optimize client identification and authentication to reduce risk and fraud	Trading algorithms that intelligently minimize risk, market impact, fees, etc.	Attempt to observe unexpected relationships between securities and indicators	Streamline, optimize, and automate tedious functions to better focus human effort	Risk mining increases risk transparency by removing filtering and aggregating layers of human reporting.
NLG	Chat functions for on-demand queries of clients and employees		Regularly generate up-to-date portfolio reports and risk commentary	Generate regular and on-demand MI reports, with language and detail adjusted to the target audience	

useful even in critical situations that are customer facing and/or subject to regulatory oversight.

Table 1 shows a number of use-cases in asset and investment management. One way to characterize the use-cases is given by their positioning in the value chain of asset management. We distinguish between:

- Client facing functions.
- External, market-oriented perspectives.
- The internal view on the actual portfolio management.
- Opportunities in the supporting functions in the middle- and back-offices.
- The risk management view.

The other dimension is segmented by the types of AI-based algorithms that are used to perform a given use-case. In order to avoid talking about "AI" in too generic terms, we differentiate the field into:

- AI/ML algorithms for pattern recognition, where learning can happen supervised, unsupervised, with reinforcement learning, and more recently with transfer learning and synthetic data.
- Natural language processing (NLP), with a focus on analyzing and extracting sentiments or intents.
- Natural language generation (NLG), with a focus on automated production of texts that are intended for human consumption.
- More recently, emerging intentions to apply behavioral analysis to emerging trends from data (not only finding trends, but also answering why they are there)

In this collection of use-cases, we can find some applications of AI and ML that are already proven in practice and reliably creating benefits. These are mainly in the middle/back office functions, where automation and streamlining based on machine learning have an immediate (cost and/or risk reducing) effect on the amount of remaining human efforts to be deployed.

The language-oriented use-cases for NLP and NLG are still constantly evolving, and there is ample room for improvement as different challenges arise in different languages (grammar, meanings, synonyms). In specific, customer facing applications need to deal with clients that in some cases mix multiple languages into their communications with the banks' chat bots or other channels (especially, the multicultural locations like Hong Kong, Singapore, among others, experience these challenges). More straightforward are the use-cases in which the AI is processing written (corporate) reports and news statements. Blackrock stated that their algorithms to detect signals in earning guidance are analyzing 5,000 earnings call transcripts per quarter and more than 6,000 broker reports every single day [Blackrock (2019)]. Extracting relevant explicit information, as well as other implicit sentiments from those texts is aiming to mimic the basic first level activity that a human analyst would do. With this, the machine is able to cover a wider scope of texts and to summarize the contents for the next level of activity.

> *"Applying AI (from creation, training, up to interpreting intelligent systems) requires a good deal of skills from the human collaborators."*

Common applications of AI algorithms are the attempts to authenticate users (to reduce risks, fraud, anti-money-laundering cases) or to identify trading patterns for individual participants or, e.g., industry groups. Optimized trading algorithms are frequently implemented, where we can state (in accordance to the incentive alignment argument of section 1) that there are different performance indicators to be optimized. In some cases, the predictive pricing has priority and the AI aims to achieve best returns for the trades. In other cases, the optimization objective can be to reduce market impact (on market prices) and trading risk or trading fees. Similarly, minimization of margin requirements, thus optimizing a banks' regulatory capital via AI-based control of risk-weighted assets (RWA) and better margin valuation adjustments (MVA) [FSB (2017)].

More recently, the discipline of behavioral analytics is emerging stronger as it appears that pure data analytics alone often does not capture the actual intentions, concerns, and incentives of the actors in the markets. The models developed based on factual data are being enhanced by the behavioral analysis and assumptions on human market participants. So-called nudges are one way to influence users and clients to either start to think about financial options or to reconsider their decisions when they appear to be non-rationally skewed towards a non-beneficial outcome [see Suh (2019) for some thoughts on AI and nudges].

Another rather new idea is the implementation of AI into corporate wide "risk mining" activities. Risk Mining is the idea to enhance the existing risk management framework in a company by a real time, interactive component that triggers individual employees to think about and report on risks. For any kind of risk (technical, financial, reputational, cultural, and the like) a risk catalogue is provided to break down risks into risk questions that help to analyze and inquire risks in more detail. Those risk questions are provided to staff via their favorite communications channels (messengers on phones, desktop systems, etc.), rather frequently but in small, acceptable doses. The objectives of the generated risk questions and their answers arise either from company-specific scenarios (which might be reported first in other departments) or from industry-wide developments (of regional/global scale, seen and felt by many or all companies in a certain industry). Result of the risk mining is a real-time management information that is not conventionally aggregated, across the levels of corporate hierarchy. Thus, real-time insights can be extracted that do not suffer from the usual time lag and information gaps.

There are more use-cases presented in Table 1 though, we do not explain them all in detail. This overview is aimed to provide a first level of information on the potential application settings for AI in the business areas of finance and asset management. In all cases, especially in a financial context, questions regarding responsibility, accountability, and controllability arise. Some of these we will tackle next.

3. WHAT IS OUR ROLE AS HUMANS WHEN AI IS BEING DESIGNED AND IMPLEMENTED?

AI versus humans, or rather collaborative co-existence?

From autonomous cars to sustainable investment management, there are a number of scenarios where we expect the technology to adhere to certain values that we uphold as humans in similar circumstances. We do not dive into the rather philosophical issues of decision-making in cars when it comes to life or death situations (of the people inside the car or the ones outside). We will rather focus our discussion more on the ethical aspects of applying AI and the question of what role we humans take for us, next to the AI we create.

The E.U. commission has set up a high-level expert group on AI (AI HLEG) to work out the implications of our expectations towards trustworthy AI. Acknowledging that new capabilities come with additional risks, the HLEG provided the following requirements for the creation and application of AI [AI HLEG (2019)]:

- Human agency and oversight
- Technical robustness and safety
- Privacy and data governance
- Transparency
- Diversity, non-discrimination, and fairness
- Environmental and societal well-being
- Accountability

Each of these requirements appears to be sensible and realistic in itself. In particular, in an asset management context with the required safety and security, we can identify a certain common theme across them: the principle of trust. What, in comparison, makes us trust a human advisor? It is the assumption that the advisor is well educated for the task we ask for (as measured by educational and institutional standards), reliable and consistent, valuing our privacy, acting with our best interest and fairness in mind, and able to be held accountable for the given advice.

We should expect no more and no less from an AI system. An increasing number of countries are introducing regulatory director/manager responsibility frameworks that aim to link a personal responsibility (and liability) to the management of a financial institution [Zetzsche et al. (2020)]. Applying AI for tasks in the financial value chain must not defer the responsibility away from the human. Thus, human agency and oversight will give us the right (and requirement) to ask the questions of the system and receive solution options or recommendations. We humans will be supported by the system, in our information collection, reasoning, and decision-making. However, we keep oversight over the process by mechanisms like the "human-in-the-loop" (HITL) principle, which enables human insight and intervention during all stages of AI system activity. Tang (2020) proposes a framework of six HITL paradigms that help differentiate the ways humans and AI systems interact (Figure 1).

The six paradigms capture the various roles and different skill sets that are required to efficiently work with AI. The AI system creator, for example, needs a completely different skill set that of an AI user, but most importantly, even the AI user should be educated in the basics about the type of AI algorithm at hand, its strengths and weaknesses, its scope of application,

Figure 1: Human-in-the-loop paradigms to ensure human agency

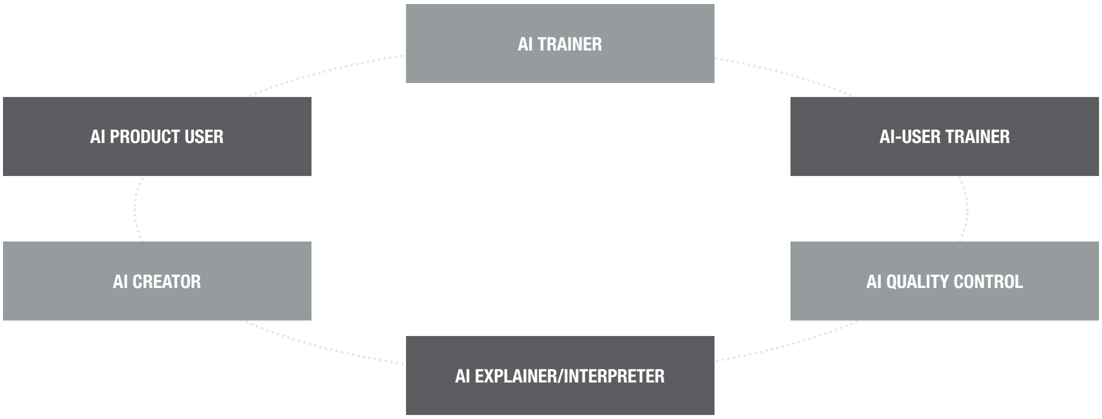

and the proper way of training and interpreting the results. AI quality control is often executed by internal corporate boards that oversee the alignment of AI and AI strategies with the mission, vision, and purpose of a company.

Providing opportunities for staff to acquire the right skills for their intended role in the AI context is the first and foremost step when embarking onto the AI journey.

4. TOKENIZATION OF ASSETS AND THE CHALLENGES OF REAL-TIME SETTLEMENT

Convergence of AI and blockchain technologies

Now, returning to more financial technology, we want to introduce a use-case for AI that is not explicitly mentioned in Table 1, is entering the marketplace from a non-conventional angle, and has the power to disrupt the ways of trading and investing as we know them today. The next big movement in asset management will likely be the currently evolving tokenization of "everything", especially of higher value private assets. By tokenization, we mean the creation of a digital representation of an asset and that this representation typically can be easily fractionalized in a simple and scalable way. Digital representations of assets can be held on a blockchain, as a form of distributed ledger technology (DLT). A distributed ledger in finance and asset management is largely an immutable journal of ownership transfers (transactions) that are held private to the participants in these transactions. Transactions and resulting balances of participants are agreed upon by a specific consensus mechanism (of which there are several options, depending on the particular setting of the blockchain).

Blockchain technology came to fame with the advent of Bitcoin [Satoshi (2008)] and other crypto currencies, which are distributed via an unpermissioned network, meaning that there is no control and restriction regarding who is participating anonymously in the network. Applications in financial institutions require, however, that the participants be subject to a "know your customer" (KYC) process, so anonymous participation is ruled out. It also makes the computation of consensus more cost effective, if the number and type of participants is limited to trusted parties in the DLT network. Many examples for enterprise grade blockchain-applications are developed on the R3 Corda framework [Brown (2018)].

Issuance of tokenized assets is then possible in two different ways: (i) issuance of "asset-backed" tokens, i.e., there exists a regular asset, often certified on paper, and the digital tokens are merely an electronic pointer to (fractions of) this certificate, or (ii) issuance of natively digital tokens, for which no other representation exists than the digital token itself, typically issued on an immutable ledger on a DLT. This DLT can be run by the issuer or by an intermediary that offers additional services like key storage and management (instead of asset custody in the non-digital case).

In both ways of issuance, the result is a token representation of an asset that will be tradeable instantly through triggering the change of ownership on the DLT. In the standard settlement of paper certificates, the monetary payment can typically be executed faster than the settlement of securities under custody, which happens via the different intermediaries in the custody chain. Contrary to that, since token ownership is transferred instantly, the DLT-settlement is requiring an equally instant settlement of the purchase amount. This in turn requires a form of digital money available for the "delivery-versus-payment" (DvP). In the token space, where the digital asset is exchanged instantly against digital money, we can rather talk about "token-versus-token" (TvT) instead of DvP.

Those payment tokens can be utility coins (like the JP Morgan coin) or stable coins or other platform/specific tokenized versions of cash that allow for instant exchange. The long awaited, prepared (and sometimes feared) advent of "central bank digital currencies" (CBDC) will offer the electronic fiat version of digital payment on DLT. At the time of writing of this article, China and a few other countries are expected to be close to the introduction of the digital version of their sovereign currency.

Risk mitigating functions like the central counterparty (CCP) are not required anymore, if instant settlement of the TvT process is achieved. But in order to provide such a risk-free settlement, the system must provide settlement finality, i.e., guarantee the irreversibility and irrevocability of the DLT transaction.

Assuming the DLT realization is able to provide instant settlement of the tokens and the digital money, a new challenge arises that can be approached with the help of AI: the immediate provisioning of this cash from a treasury perspective. Today, the treasury department of a market participant benefits from the T+1 or T+2 settlement duration by netting the trades during the day and being able to provision only exceeding amounts for payment to counterparties on the following day(s). This will change fundamentally with TvT. As each trade is settled in real-time, netting up buys and sells will not be possible anymore. ISSA reports that, for example, a daily gross settlement obligation volume of U.S.$1.3 trillion today are netted to only U.S.$19.8 billion and collaterized by U.S.$7.3 billion in margin deposits [ISSA (2019)].

While smaller scale "daily business" arguably leads to a largely balanced stream of TvT of similar size on buy side and sell side, the required buffer in digital cash is likely not extremely big compared to the trading volume of a market participant. Also, even in today's non-digital process, market participants host sizeable volumes of cash on their depositary accounts with the central banks. Retail business might be balanced, as consumers buy and sell in similar (uncritical) sizes under different market conditions. Critical situations will arise in commercial and institutional segments, when large funds and other investors might create a large-scale unbalanced movement on the buy side. Such a purchase must immediately be supported by availability of tokenized cash, pre-funded on the TvT accounts. This has limiting implications on the liquidity and cost, as well as on the possibility of market makers to provide intraday liquidity and trades to the market.

One option to tackle this challenge is the attempt to predict in specific the large scale of movements that are not covered by the average pre-funded amounts. Even short-term reaction times, gained by an algorithm that anticipates movements shortly before they happen, would be helpful. Treasury will then use digital exchanges to provide for more tokenized cash generated from non-digital reserves and be guided in this process by short term prediction systems.

Figure 2: Schematic system diagram for a one-day predictive solution

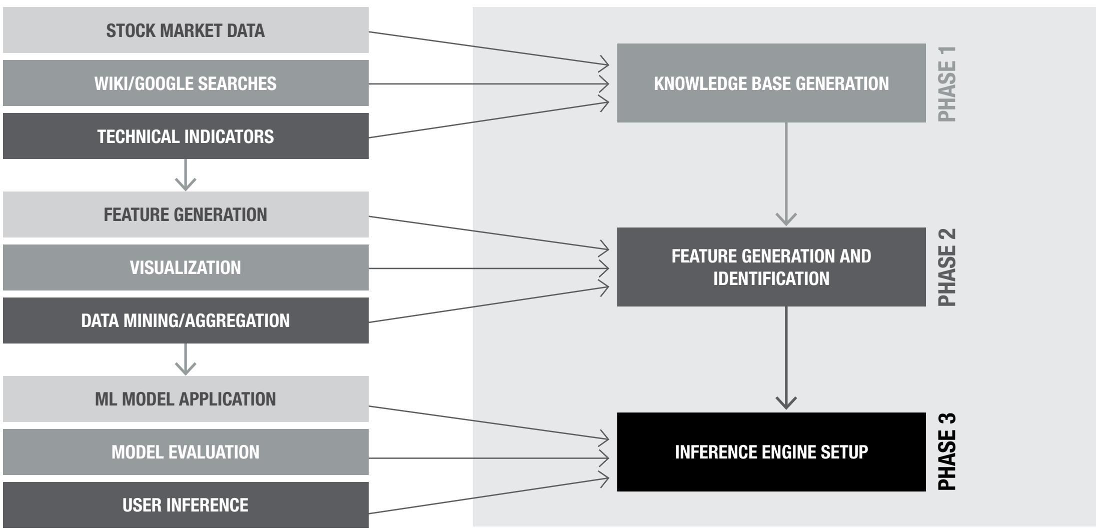

Source: Weng et al. (2017)

Weng et al. (2017) have proposed an approach for such an AI prediction engine that is not only based on historical time series data. It also takes into account data from very current news feeds and information from crowd-sourced inputs (e.g., social media) as well as trends in user visits or query statistics from search engines or information databases (Google, Wikipedia). The principal objective of such a system is not to excel in the precise prediction of a large set of stocks and their future prices. The important quality is to use the variety of inputs as sensors for surprising, potentially large movements in an asset, which would require an unprepared amount of liquidity in a short period of time. Figure 2 shows the basic functionality of this prediction system, structured into processing phases, with the more detailed sub-functions that establish a good current view into the market and its predictive indicators.

Another upcoming challenge with tokenized assets lies in the fact that – especially in the early years – several assets might exist in an on-chain version (i.e., DLT-based) as well as in an off-chain version. For the investor community, this reduces the perceived risk with this rather new class of tokenized assets, as the 1:1 exchangeability of digital and non-digital version would be guaranteed by their financial institution to safeguard the holdings in any case. However, it can be expected that there is a difference in trading volume and trading speed, as well as liquidity, in such competing markets, resulting in (slightly) differing prices due to the bifurcation of liquidity into both markets [OECD (2020)]. Such arbitrage opportunities will attract market participants who try to exploit these differences by "intelligent" trading algorithms, or even create such opportunities by a set of actors on both markets, who execute strategies to drive prices on one market and harvest the benefits on the other.

5. CONCLUSION AND OUTLOOK

The only winning move is to skillfully play

In this article, it has been shown that AI in finance and asset management has been implemented in a large variety of use-cases already, with constantly more emerging across the value chain.

The major take-away is the recommendation not to view AI as a simple plug-and-play tool. Applying AI (from creation, training, up to interpreting intelligent systems) requires a good deal of skills from the human collaborators in order to harvest the benefits for the intended use-case.

Applying AI to asset management must be seen as part of the more general digitalization journey. The challenges are, therefore, not only in the understanding of the technology. Major roadblocks can be the conservative collective mindset in the corporation, avoiding new opportunities, or the reluctance of the top management to invest time, money, and education in the workforce that will be tasked to "work on" and "work with" new technologies. As ever so often: new skills help to understand, to implement, and to apply new technology. Not to play is not an option.

REFERENCES

AI HLEG, 2019, "Ethics guidelines for trustworthy AI," Artificial Intelligence High level Expert Group, European Commission, April

Blackrock, 2019, "Artificial intelligence and machine learning in asset management," Blackrock Public Policy View Point, October, https://bit.ly/2TthLtu

Brown, R. G., 2018, "The Corda platform: an introduction," R3 Corda documentation, May, https://bit.ly/2vwceKC

FSB, 2017, "Artificial intelligence and machine learning in financial services," Financial Stability Board, November, https://bit.ly/2VD2toH

Google, 2019, "Perspectives on issues in AI governance," Google, https://bit.ly/2x0QFCd

ISSA, 2019, "Crypto assets: moving from theory to practice," International Securities Services Association, November, https://bit.ly/38aBkvO

Minsky, M., 1974, "A framework for representing knowledge," MIT-AI Laboratory Memo 306, June, https://bit.ly/2PGhO3Y

Murphy, T., 2013, "The first level of Super Mario Bros. is easy with lexicographic orderings and time travel . . . after that it gets a little tricky," Carnegie Mellon University, https://bit.ly/3ak46vn

OECD, 2020, "The tokenization of assets and potential implications for financial markets. OECD Blockchain Policy Series, January, https://bit.ly/2Thz3e5

Satoshi, N., 2008, "Bitcoin: a peer-to-peer electronic cash system," https://bit.ly/3alXFYR

Suh, R., 2019, "Can AI nudge us to make better choices?" Harvard Business Review, Behavioral Economics, May

Tang, B. W., 2020, "The Chiron imperative – a framework of six human-in-the-loop paradigms to create wise and just AI-human centaurs," in Bhatti, S. A., S. Chishti, A. Datoo, and D. Indjic (eds.), The LegalTech book: the legal technology handbook for investors, entrepreneurs and fintech visionaries, Wiley, forthcoming

Vontobel, 2018, "Artificial intelligence in investing. Great aspirations – reverie or reality?" Vontobel Asset Management, May, https://bit.ly/2vzUQEP

Weng, B., M. A. Ahmed, and F. A. Megahed, 2017, "Stock market one-day ahead movement prediction using disparate data sources," Expert Systems with Applications 79,153–163

Zetzsche, D. A., D. Arner, R. Buckley, and B. W. Tang, 2020, "Artificial intelligence in finance: putting the human in the loop," CFTE Academic Paper Series 1, https://bit.ly/3cqAI8P