CAPCO

a wipro company

THE CAPCO INSTITUTE
JOURNAL
OF FINANCIAL TRANSFORMATION

ORGANIZATIONAL

The truth behind artificial intelligence: Illustrated by designing an investment advice solution

CLAUDE DIDERICH

ARTIFICIAL INTELLIGENCE

#58 NOVEMBER 2023

THE CAPCO INSTITUTE

JOURNAL OF FINANCIAL TRANSFORMATION

RECIPIENT OF THE APEX AWARD FOR PUBLICATION EXCELLENCE

THE TRUTH BEHIND ARTIFICIAL INTELLIGENCE: ILLUSTRATED BY DESIGNING AN INVESTMENT ADVICE SOLUTION

CLAUDE DIDERICH | Managing Director, innovate.d[1]

ABSTRACT

Artificial intelligence can be considered one of those technologies, like 5G, 3D printing, and virtual reality, that can disrupt the business world. While AI has the potential to solve meaningful business problems, implementing it in a way that creates value is challenging. Unfortunately, many AI proponents lack the necessary computer science and mathematics machine learning skills required for developing AI systems that pass the Turing test. This paper presents an assessment of the characteristics of AI, allowing the reader to understand what specific business problems it can solve, and describes how an AI-supported investment advice solution for wealthy private clients can successfully deliver value. By reviewing the lessons learned, I conclude that the future of AI is bright if the focus is put on applying it to those challenges that it is best suited to solve.

1. INTRODUCTION

Artificial intelligence (AI) can be considered as the most disruptive technology dominating the 21st century [Girasa (2020), Roubini (2022)]. While the concept of AI may sound frightening, it is impossible to ignore the value that it offers society. Semi-autonomous cars (e.g., Tesla), spoken language recognition (e.g., Siri, Alexa), purchase recommendations (e.g., Amazon, Netflix), subject tracking in photo and video cameras (e.g., in the latest products by Canon, Nikon, and Sony), as well as chatbots (e.g., Bard, BioGPT, and ChatGPT), are just some examples of solutions that rely heavily on AI. In addition, though not widely recognized, AI supports numerous non-end-user facing activities, like detecting possible credit card frauds, pricing insurance risks, or constructing investment portfolios.

Undoubtedly, AI has already destroyed certain jobs and will continue to do so in the future. However, it will also lead to the creation of new jobs, help humans focus on more creative activities and improve their lives. It will allow firms to increase efficiency, resulting in lower costs and prices, extend product customization, allow customers to solve their problems in better ways, and extract meaningful data patterns from big, complex datasets, supporting superior decision making. AI offers customers a vast treasure of value-creation capabilities and value-appropriation opportunities for organizations [Kaartemo and Helkkula (2018), Wodecki (2018)].

Despite the aforementioned potential benefits, human judgment will continue to be needed to identify those challenges that can best be solved by using AI and to design the respective solutions. Human judgment is required to determine which problems AI can and cannot solve and what data AI needs to learn and create new insights from, how AI can supplement human intelligence and where it can replace it, as well as addressing the ethical challenges associated with relegating recommendations and decision making to "digital humans" [(Diderich (1993), Chancellor (2023)].

[1] The author acknowledges the valuable feedback from Esther Gelle on an earlier version of this paper.

1.1 Key moments in the history of AI

Taking a short excursion into the history of AI helps us better understand AI and its value to society. While AI may be seen as a recent phenomenon, it finds its roots in the 1950s. British polymath Alan Turing first suggested that if humans can use information and reasoning to solve problems and make decisions, computers should be also able to do so. This led him to formulate the famous Turing test. Originally called the imitation game [Turing (1950)], it assesses a machine's ability to exhibit intelligent behavior equivalent to, and indistinguishable from, that of a human. As of today, no general-purpose AI system has passed the Turing test.

AI started flourishing in the late 1960s as computers became more accessible. Early work by Feldman, Feigenbaum, Minsky, Newell, Simon, Weizenbaum, and Winograd [Barr and Feigenbaum (1981)] showed promising results in applying goal-based problem solving using expert systems. Expert systems [Puppe (1993)] focus on encoding rules of human thinking into computer programs. They look literally like sophisticated "if-then-else" programs. Specialized computer languages, like Prolog and Lisp, were developed to support encoding human decision rules efficiently and effectively. Designed to focus on specific problems, expert systems were high-performing, transparent, very reliable, and offered easy-to-understand results. They were well-suited for targeted problem solving. Some of the most prominent expert systems in medicine were MYCIN (diagnosing and treating infectious diseases), DENDRAL (molecular structure prediction in chemical analysis), and CaDet (detecting cancer in early stages). All these systems are based on modeling how humans understand their decision-making process rather than how human brains work. Their main drawbacks are that they lack generality and require extensive maintenance work to update the rules.

In the 1980s, AI got a second lease of life when computing power allowed for general-purpose "artificial neural networks" (ANNs) to be trained and deployed at scale. ANNs are based on the structural understanding of the human brain rather than on encoding human decision making. Haykin (1999) describes a neural network as a "massively parallel distributed processor made up of simple processing units – the neurons, which has a natural propensity for storing experiential knowledge and making it available for use. Knowledge is acquired by the network from its environment through a learning process. Interneuron connection strengths, known as synaptic weights, are used to store the acquired knowledge." Its simplest form is the perceptron, a linear classifier [Rosenblatt (1958)]. ANNs are generic, problem-independent functions that use previously learned insights to calculate solutions based on input data. The insights ANNs use are derived using deep learning algorithms on typically large sets of historical problem-solution data rather than hard-coded algorithms. Work by, among others, Fukushima, Grossberg, Hopfield, Kohonen, Linkster, Peral, Schmidhuber, Rosenblatt, Vapnik, Waibel, and Werbos, substantially advanced the field of knowledge in machine learning, which stands at the center of today's AI systems [Haykin (1999)].

An important, but often forgotten, result by Amaldi and Kahn (1995; 1998) and Engel (2001) finds that optimally training a single perceptron, the simplest possible ANN, is an NP-hard problem.[2] They show that it is computationally intractable (unless using quantum computers) to train an ANN to always produce the correct solution. Moreover, it is even impossible for an ANN to be trained in a way that it finds the best possible solution within a given degree of certainty; that is, in a probabilistic sense.

Few advancements were made in AI during the 1990s. However, in the new millennium, research in machine learning got a new boost with the advent of high-speed computers, low-latency networks, and massive storage capacities, making collecting and processing large problem-solution datasets more tractable [Hwang (2018)]. Furthermore, the non-scientific community also became interested in AI, notably due to landmark achievements, such as IBM's DeepBlue beating Kasparov in chess in 1997 [Campbell et al. (2002)], Google's self-driving car experiment since 2009 [Teoh and Kidd (2017)], the Watson computer system winning the first prize on the "Jeopardy" quiz show in 2011 [Baker (2011)], Apple's voice recognition Siri (2011) and Amazon's Alexa (2014), AlphaGo winning multiple games against Go champions since 2015 [Pumperla and Ferguson (2019)], and more recently DeepL (2017) and ChatGPT (2023) relying on generative AI methods and making large language model technology [Zhao et al. (2023)] available on mobile devices.

[2] A problem is called "NP-hard" if it is suspected that no algorithm using polynomial time (versus exponential time) exists that can solve it.

2. UNDERSTANDING THE "INTELLIGENCE" IN ARTIFICIAL INTELLIGENCE

Consider the following business idea: an entrepreneur wants to offer customized pizzas in any shape or form, like, for example, a heart shape for Valentine's Day, a steamboat pizza for celebrating a child's birthday, or a tennis racket in honor of the 100th anniversary of the Italian Open. One challenge the entrepreneur faces is estimating the amount of tomato purée needed for each pizza. I call this problem the tomato purée pizza challenge or TPPC.

The traditional approach for solving the TPPC would be calculating the amount of tomato purée using a distinct formula for each shape. For example, $f_{square}(d) = 0.2d^2$ for a square pizza and $f_{heart}(d) = 0.2(1 + \pi/4)\ d^2/2$ as an approximation for a heart-shaped pizza. While the result of this approach is exact, transparent, and quick to calculate, it lacks flexibility. A new formula has to be developed and encoded for each new pizza form. Intelligence is associated with the different formulas $f_x()$ developed.

Addressing the TPPC using AI takes a different route. First, a large number of different pizza shapes are designed (and recorded as images). The amount of tomato purée required for each is measured empirically. This leads to an extensive problem-solution dataset *D* of pizza image-tomato purée pairs. Next, a generic ANN is trained using a supervised learning algorithm, as found in standard AI algorithm libraries, on the dataset *D*. Insights *I* are derived through learning from the dataset $D : I = \text{ANN}_{learning}(D)$. Finally, for a given image *p* of a pizza shape, the trained ANN calculates the amount of tomato purée needed using the generic function ANN(*I, p*). Intelligence is associated with the insights *I* dynamically learned from the data rather than a hard-coded formula.

In contrast with analytical approaches, the AI solution works for any pizza shape rather than only a pre-coded subset. However, it will only deliver reasonably correct results if the ANN has been trained using a representative and large dataset of pizza image-tomato purée pairs. Furthermore, the calculated amount of tomato purée required may be way off for some pizza shapes.

This paper aims to understand when AI is an appropriate tool for solving a problem and when other methods are more suited. Different problems require different solution approaches.[3] For some problems, AI will be the most appropriate approach; for others, different solutions will prevail. Consequently, when deciding whether to rely on AI to compute a solution for a given problem, it is important to:

1) Understand what the "exact problem" that needs to be solved is.
2) Know what "historical data" is available and can be legally used for learning and insights generation.
3) Know the "value and limitations" of using possible AI solutions.

3. DETERMINING THE SUITABILITY OF AI FOR SOLVING A SPECIFIC PROBLEM

While AI can be used to approach many wicked problems,[4] as solving the TPPC has shown, it is by no means applicable to solving every problem. The universe of problems most suitable for AI can be classified into two categories:

1) **Pattern-matching problems:** problems in this category are solved by identifying complex structures or patterns in datasets and associating them with specific solutions. An example of a typical pattern-matching problem is image recognition, e.g., identifying a cat or a human crossing a street in a picture. Speech recognition, matching spoken waves to words, is another such problem. Playing games like Go or chess can also be handled using AI algorithms designed for pattern matching. Recently, chatbots like Bert, ChatGPT, or Galactica have used large language model algorithms to solve generic pattern-matching problems, matching chat questions to learned text. Problems in the pattern matching category are best addressed using "supervised" learning algorithms [Jo (2022a)] applied to labeled datasets. The term supervised relates to the requirement that the training dataset includes labels representing known solutions to specified problems.
2) **Classification problems:** the second category of problems well suited for AI algorithms are solved by classifying data based on unknown attributes. AI can address typical classification problems: customer segmentation, anomaly detection, or product recommendations. Unsupervised learning algorithms are typically used to address classification problems [Jo (2022b)]. In contrast with analytical approaches to solving

[3] Note for the sake of completeness that in the case of the TPPC, analytical algorithms exist, for example, using triangulation, which can approximate the surface of a generic shape without relying on AI.

[4] A wicked problem is a problem that is difficult or impossible to solve because of incomplete, contradictory, and changing requirements that are often difficult to recognize.

classification problems, these algorithms do not need to know a priori what attributes are relevant for classifying the data. During unsupervised learning, the relevant attributes are determined implicitly and often remain hidden from the outside world. Unlike supervised learning algorithms, unsupervised learning does not require solution data; that is, training data can remain unlabeled.

Identifying what value can be created by, and appropriated from, using AI as a problem-solving approach is critical. AI has the potential to deliver significant value in two business areas:

1) Identifying patterns or attributes that are "too complex" or "take too much time for humans to identify", especially because of their multi-dimensional nature or the size of the dataset. Typical problems in this category are constructing investment portfolios, detecting credit card fraud, or tracking image data.

2) "Performing repetitive tasks", where AI is significantly faster and/or cheaper than human resources. Typical problems in this category are voice recognition, text translation, writing draft documents, or searching for specific data items.

3.1 Five premises for using AI-based problem solving

Not all problems are sound for AI solving. Five premises must be satisfied to solve a wicked problem using AI successfully. These are:

(1) The problem at hand "cannot be solved using analytical algorithms", or using analytical algorithms is computationally infeasible, although theoretically possible.

(2) The problem "can be solved by relying on available historical data". The solution is not entirely novel. This does not mean that existing data must include the solution but that it can be reasonably inferred from it.

(3) There exists appropriate "labeled" or "unlabeled datasets" (depending on the type of problem) that can be legally used for AI learning purposes.

(4) Relying on a suboptimal or incorrect solution "is a viable option".

(5) The problem solver is "not faced with any moral hazard due to an incorrect solution" computed by AI. No human lives are at risk if AI fails to find the right solution.

Premise (1) states that a problem-specific algorithm exhibiting validated properties is preferred to a problem-agnostic machine learning approach. Although this may seem obvious, it means that AI should not be used as a replacement for human domain-specific knowledge. While AI allows for combining existing knowledge in a way that humans might not have thought of, premise (2) states that AI cannot invent new knowledge. With efficient data-collecting resources available, satisfying premise (3) should be straightforward. However, it is not. Legally collecting high-quality data often poses an insurmountable challenge. Finally, premises (4) and (5) address the challenge that AI cannot guarantee the correctness of its results. In many situations, AI cannot even offer the reasoning that has led to the solution, thus making the work of human result validation tedious, if not impossible. Recent research in explainable AI [Holzinger et al. (2022)] focuses on addressing that challenge. Premise (5) stipulates that if the use of AI could lead to moral hazard, it must be used primarily as a decision-support tool complemented by human expertise and/or analytical algorithms.

3.2 Challenges faced by AI

Applied to the right problems, AI can offer solutions humans could not think of. However, these solutions have some caveats that must be understood before relying on them.

Whether relying on labeled or unlabeled data for learning, generic AI algorithms make it possible to find correlations in the training data, but not causalities. Pearl (2000) and Pearl and Mackenzie (2018) have shown that "data alone can never answer causal questions. They [AI algorithm developers] require to formulate a model of the process that generates the data or at least some aspects of that process." Incorrectly assuming causality when only a correlation exists is one of the biggest mistakes one can make when relying exclusively on data to solve problems. This is no different for AI. Many, if not all, sophisticated AI algorithms include some sort of domain-specific model to support deriving causalities. For example, ChatGPT has learned that most famous sports journalists have covered the Olympic games. However, when ChatGPT is asked what events a known sports journalist has covered, it incorrectly infers the causality that such a journalist must have covered, with a high probability, the Olympic games, although only a correlation exists.

Training a generic ANN in such a way that it correctly classifies the largest possible number of data elements is a computationally intractable problem. This means that unless using domain-specific modeling when designing and training an ANN, it is impossible to ensure, even in-sample, the quality of any result. It is computationally infeasible to train a generic ANN in such a way that it correctly solves a given percentage of problem instances; that is, offers a probabilistically correct answer.

A third, and even more vital, challenge that many AI systems face is that they are black boxes. AI typically provides a possible solution but cannot explain how that solution was derived. For example, an ANN used for recognizing animals in images was trained using, among others, horse images that included a copyright notice (which non-horse images did not have). When using the trained ANN on new images, it incorrectly identified any image containing a copyright notice as an image of a horse [Lapuschkin et al. (2019)]. Although research in designing explainable AI (XAI) algorithms has made progress in recent years [Samek et al. (2017)], notably by attributing the statistical probabilities of each input to the result component, there is still a long way to go to come up with domain-independent AI algorithms that offer explainable solutions. Most promising research in XAI focuses on designing interpretable models using decision trees, Bayesian networks, and sparse linear models [Rudin (2019)].

4. TOWARDS MORE SUCCESSFUL INVESTMENT ADVICE: AN AI CASE STUDY[5]

Offering "customized investment advice" (CIA) as a paid service has become one of the most prominent offerings in private banking. One of the reasons for this is tighter regulations imposed on investment advisors to avoid conflicts of interest. Another is that customers seek help navigating the ever more complex investment universe without delegating the final investment decision. Private banks like CIA because it can be sold in a way to generate recurring revenues.

4.1 Understanding CIA

A naïve manager would consider CIA a product recommendation problem, similar to Amazon suggesting to its customers which books to buy based on past purchases or Netflix proposing what movie to watch next based on learned user preferences. Unfortunately, advising CIA customers is more complex, as it involves multiple stakeholders with different goals and preferences: the investor as the customer, the investment advisor and their employer as service providers, and the investment product providers as the manufacturers.

Investors look for investment recommendations that meet their risk profiles, reflect their market expectations, and bring them closer to their financial goals. Investment advisors aspire to advise clients effectively, maximizing the probability that the investor will act upon their advice and be happy. They also look for help navigating the ever-growing universe of investment products, each with its features and caveats. Their employer, on the other hand, wants to maximize value capturing. Finally, investment product providers look for their offerings to be recommended by the investment advisor. Based on these observations, the CIA service can be reformulated as a decision problem suitable for solving using AI.

4.2 Formulating the CIA as an AI problem

To determine which products to recommend to its client, the investment advisor must evaluate the function shown in equation (1), where the parameters ❶, ❷, ❸, and ❹ are defined in Figure 1 as the data universe used in offering CIA.

$$r(❶, ❹, a(❶, ❷, ❸, ❹)) = \{\text{set of advised investment products}\} \quad (1)$$

Function $r()$ is an analytical function encoding applicable regulations ensuring that investment recommendations align with them. The $a()$ function computes the set of products from the product shelf that the investment manager recommends to the investor based on their investment goal, risk profile, subjective market views, preferences, and given objective market conditions. While function $r()$ fails the AI premises (1), (4), and (5), defined in Section 3.1, function $a()$ meets all five AI premises. As such, function $a()$ is well suited to be implemented using a multi-layer feed-forward ANN.

[5] Note that the case study presents a high-level description of how AI can be used to improve the customized investment advice service. For the sake of readability, the description has been simplified. Details, including some modeling and intermediary data processing steps, have been omitted.

Figure 1: Data universe used in offering CIA

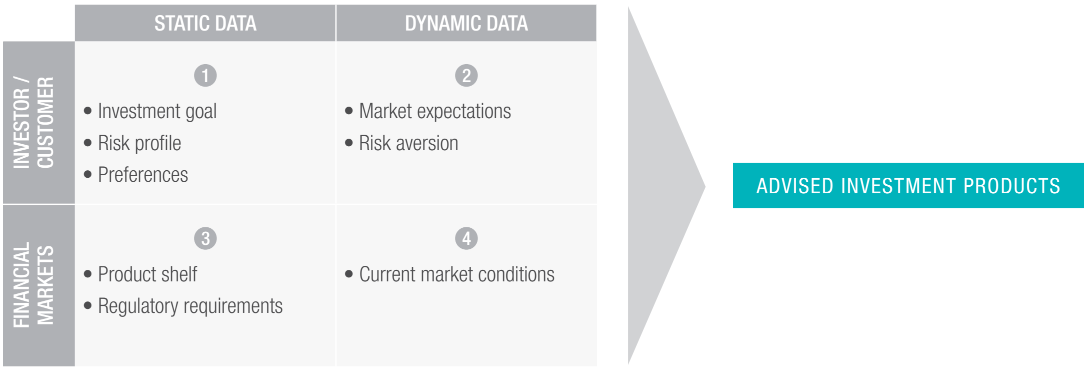

Notes: Static data is time-independent, whereas dynamic data changes over time.
Investor/customer data is specific for each investor, whereas financial market data is the same for all investors.

Table 1: AI technology stack model for implementing and operating an AI solution

TECHNOLOGY STACK	DEFINITION	FUNCTIONALITY FOCUSING ON THE CIA PROBLEM
Infrastructure platform	Hardware underlying the AI solution	Generic AI implemented on a cloud infrastructure (i.e., AI as a service)
Framework	AI architecture	Multi-layer feed-forward ANN
Learning algorithm	Specific machine learning algorithm(s) used	Off-the-shelf, supervised learning algorithm
Data pipeline	Data source and data management platform	Proprietary client data, proprietary market and risks assessment data, public market data
AI service	Well-defined service applying the learning algorithm to the data pipeline, consistent with the framework using the infrastructure platform	General purpose API makes it possible to learn/calculate the parameters; that is, the weights associated with the nodes, of the multi-layer feed-forward ANN
Scoring algorithm	Domain-specific AI solution addressing the business problem	Custom-build capabilities resulting in investment product recommendations based on client-specific data and current financial market conditions

Source: Based on Tsaih et al. (2023)

To develop and train an ANN that implements the function $a()$, I use a variation of the AI technology stack proposed by Tsaih et al. (2023), as shown in Table 1. A structured approach focusing on specific outcomes in a well-defined order helps avoid mixing different concepts, which could result in a sub-optimal, often even non-working, AI solution. It also helps make it easier to identify the exact problem that needs to be solved by distinguishing between technology, framework, training data, learning, and scoring. By encapsulating all historical data aspects into the data pipeline stack, the approach ensures that the appropriate data is available and can be legally used. Finally, distinguishing between learning and scoring algorithms helps identify value and limitations of the AI designed solution.

4.2.1 INFRASTRUCTURE PLATFORM

As the problem to be solved is a typical, although multi-stakeholder, pattern-matching problem, there is no need for a problem-specific AI infrastructure platform. Furthermore, due to the difficulty of estimating the computing resources required for training and scoring the ANN a priori, I rely on a generic AI cloud infrastructure such as Amazon AWS AI, Microsoft Azure AI, or IBM Watson ML.

4.2.2 FRAMEWORK

Next, I model the function *a*() as a multi-layer feed-forward ANN, where the input layer ingests the parameters ①, ②, ③, and ④, excluding regulatory requirements. The output layer is associated with the recommended investment products from the product shelf.

To keep the designed framework as simple as possible, I refrain from integrating back-propagation that would allow the ANN to learn by itself from the market performance of the recommended investment products while scoring and correcting faults in internal stages of the network. Instead, I regularly re-train the ANN when relevant new investor and market data becomes available.

4.2.3 LEARNING ALGORITHM

While designing the third level of the AI technology stack, I use a standard supervised learning algorithm offered by the cloud infrastructure platform rather than developing a proprietary one. Such algorithms typically depend on gradient-driven optimization combined with heuristics to speed up the computations and avoid local optima.

4.2.4 DATA PIPELINE

Organizing and managing the data pipeline is the most challenging part of designing, building, and implementing an ANN. Each grey vertical box in Figure 2 (training data) represents a separate dataset for training the ANN. It is specific for a given investor at a given point in time. The parameters ①, ②, ③, and ④ represent the input dataset for the point in time *t*. The "investment advice" represents the output data or label associated with the input data; that is, the portfolio holdings and investment products the investor chose at time *t* as their preferred investments.

I use raw data collected from the KYC[6] process and from risk profiling the investor, as required by regulations, as static investor data (①), describing their financial goals, risk profile, and preferences. Relying on raw data allows the ANN learning algorithm to potentially identify hidden correlations between attributes while remaining fully aware of potential noise in the collected data that could negatively impact the outcome [Kahneman et al. (2021)].

Unfortunately, investor market expectations and risk aversion data are typically unavailable at a given time *t*. Consequently, I derive the investor's expectations and risk aversion (②) from their portfolio holdings at time *t*. To do so, I associate specific market expectations and risk preferences with each investment product. For example, holding technology stocks is associated with the expectation that equity markets will grow more than GDP and have a low risk aversion, whereas holding inflation-linked bonds is associated with the investor expecting inflation to rise faster than markets expect and being risk averse by seeking protection.

The third data category represents the dynamic market data (③). It describes the observed current market conditions, like inflation rate, GDP, unemployment rate, and stock market valuations. In contrast with input ① and ②, the current market conditions data is independent of any specific investor and thus identical in all datasets for a given time *t*.

Parameter ④ represents the shelf of investment products available at time *t*. Furthermore, I assume that the regulatory requirements (part of parameter ④) are codified in an analytical function and are not derived from the dataset used for training the ANN.

To label the output or learning datasets (investment advice), I assume that the investors' portfolio holdings at any given time reflect their actual investment decision. They represent the investment products that the investment advisor should have recommended to the investor at that point in time.

4.2.5 AI SERVICE

The AI service layer implements the supervised learning algorithm. It is applied to the data pipeline using the specific API the infrastructure platform provides. The outcome is a well-defined function *a*( ), which can subsequently be used to compute a set of possible investment product recommendations based on static and dynamic investor data combined with current market conditions and available product shelf.

4.2.6 SCORING ALGORITHM

Finally, the scoring algorithm calculates what investment products the investment advisor should recommend to the investor based on the insights learned by the ANN. It implements evaluating equation (1) and is illustrated in Figure 2 by the green vertical box (scoring).

[6] KYC = Know your client.

Figure 2: Data pipeline used for learning (datasets 1 to 5) and scoring (input/output)

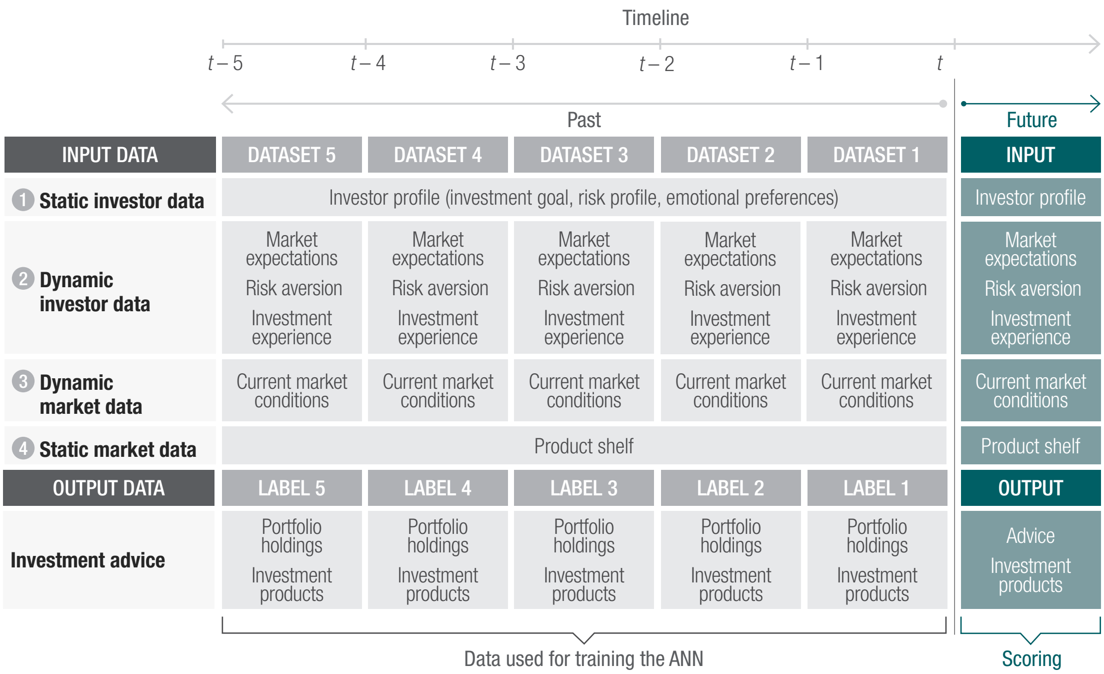

The static investor data (parameter 1) and the current product shelf (parameter 4) are extended by the investor formulating their expectations of the financial markets and their current risk aversion (parameter 2), i.e., their expectations of inflation, economic growth, future unemployment rates, etc. Current market conditions (parameter 3) are based on observed inflation, GDP, unemployment rate, etc. It is important to note that in this specific CIA solution only the investor makes predictions of the markets. The investment advisor's judgmental role is curtailed to describing the current market conditions.

Finally, the output from function *a*( ) is passed through the regulatory requirements filter *r*(), resulting in a set of investment products recommended to the investor that is regulatory compliant and available from the investment advisor's firm.

A point to remember is that the value of the investment advice depends on the model assumption that the investor was happy with their past investments made through their portfolio holdings and that they reflected the investor's market expectations.

4.3 Creating value by applying the model

The success of using AI in the designed CIA solution is based on four assumptions:

1) Similar customers (i.e., concerning investment goals, risk profile, and rationale, as well as psychological preferences) invest similarly in similar environments.

2) Investors accept that investment products advised to them may perform poorly, given their expectations.

3) Investors were happy with their past investment decisions (or the decisions that they were unhappy with were flagged as such and subsequently removed from the training dataset).

4) AI can, within reasonable boundaries, correctly identify relevant attributes in the presented datasets and classify data accordingly without requiring human modeling or manual intervention.

When these four assumptions are met, investment advice computed by AI should be expected to be superior and more consistent than guidance from human investment advisors alone. One reason is that AI has superior capabilities in identifying those attributes that matter most to investors, which may go unnoticed by human investment advisors. Another reason is that AI is able to deal better with a large universe of potential investment products than humans. Furthermore, AI does not suffer from human judgmental biases, like anchoring, availability, conjunction fallacy, optimism, loss aversion, framing, sunk costs, or overconfidence [Kahneman (2011)]. Under the assumption that the world will not be disrupted, and the future (even if non-natural and complex) will still relate to the past, using AI to support CIA should lead to a higher acceptance rate of investment advice provided and, therefore, happier customers. The potential drawback that AI cannot as yet offer explicit explanations for its advice can be mitigated by using AI as a tool to support the investment advisor's expertise rather than to replace them.

5. CONCLUSION

More is needed than just the artificial part of AI to successfully apply it to solving wicked problems. As with any problem-solving approach, only well-understood challenges can be successfully solved. One cannot expect AI to understand a poorly formulated problem, let alone solve it. This means that possessing big data is not enough. Implementing AI to create value for its users and allow its creators to capture part of that value requires diverse skills. Hard-core mathematical and computer science skills are too often left on the backbench or completely ignored.

5.1 Lessons learned from the past

Several highly relevant insights can be gained from research in AI, the presented case study, and experience implementing AI solutions to solve wicked business problems involving large datasets.

- To create and capture value in business, problem solving requires understanding the problem and identifying how a solution creates value for the stakeholders involved.
- Just because a problem involves substantial amounts of data does not make it necessarily suitable for solving using AI.
- AI is well-suited for solving problems that require identifying patterns in large datasets, which are structurally too sophisticated for the human eye to detect.
- AI best identifies correlations and correlation-like structures between data elements, especially non-linear ones, and clusters similar data elements.
- Analytical problem-solving techniques will outperform AI in most cases where computationally feasible analytical solutions to the considered problem exist.
- The most important caveat to consider when relying on AI is that it is mathematically impossible for any AI algorithm, unless combined with causality models, to guarantee the correctness of the calculated solution, even in probabilistic terms.

5.2 Looking into the future

While the lessons learned from past experiences with AI may sound grim, AI offers enormous opportunities when correctly applied. There are a considerable number of challenges where analytical approaches have failed or performed poorly. In situations where solving a problem requires mining large historical datasets, AI will outperform traditional algorithms in all but the most straightforward cases.

Two key challenges must be addressed to fully exploit AI and succeed at the Turing test. First, machine learning algorithms must include an explainable component, whether relying on supervised, unsupervised, or reinforcement learning. Black-box AI will not survive the scrutiny required for large-scale and/or mission-critical deployment. Second, AI must move from learning correlations to creating causal knowledge. As such, AI must allow for combining with human-designed causality models.

We are a long way from machines being genuinely creative; that is, creating knowledge that cannot be derived by combining existing knowledge. However, taking an optimistic-realistic approach to AI will make it possible to create and capture value beyond efficiency and effectiveness gains. AI is a sophisticated tool that, when used wisely, especially in combination with other tools (and humans), will allow for shaping critical aspects of our future.

REFERENCES

Amaldi, E., and V. Kahn, 1995, "The complexity and approximability of finding maximum feasible subsystems of linear relation," Theoretical Computer Science 147:1-2, 181-210

Amaldi, E., and V. Kahn, 1998, "On the approximability of minimizing nonzero variables or unsatisfied relations in linear systems," Theoretical Computer Science 209:1-2, 237-260

Baker, S., 2011, Final jeopardy - man vs. machine and the quest to know everything, Houghton Mifflin Harcourt

Barr, A., and E. A. Feigenbaum, 1981, The handbook of artificial intelligence (Vol. I & II), Heuristech Press

Campbell, M., A. J. Hoane Jr., and F.-h. Hsu, 2002, "Deep Blue," Artificial Intelligence 134:1-2, 57-83

Chancellor, S., 2023, "Towards practices for human-centered machine learning," Communications of the ACM 66:3, 78-85

Diderich. C., 1993, "Expert systems and their risks: an ethical study," in International Conference on Computer Science, Communication, and Society: A Technical and Cultural Challenge, Neuchâtel, Switzerland

Engel, A., 2001, "Complexity of learning in artificial neural networks," Theoretical Computer Science 265:1-2, 285-306

Girasa, R., 2020, Artificial intelligence as a disruptive technology, Palgrave Macmillan

Haykin, S., 1999, Neural networks: a comprehensive foundation, Pearson Education

Holzinger, A., A. Saranti, C. Molnar, P. Piecek, and W. Samek, 2022, "Explainable AI methods – a brief overview," International Workshop on Extending Explainable AI Beyond Deep Models and Classifiers, LNAI 13200, 13-38, Springer

Hwang, T., 2018, "Computational power and the social impact of artificial intelligence," SSRN, https://tinyurl.com/4wfytsz3

Jo, T., 2022a, "Supervised learning," in Jo, T. (ed.), Machine Learning Foundations, Springer

Jo, T., 2022b, "Unsupervised learning," in Jo, T. (ed.), Machine Learning Foundations, Springer

Kaartemo, V., and A. Helkkula, 2018, "A systematic review of artificial intelligence and robots in value co-creation: current status and future research," Journal of Creating Value 4:2, 211-228

Kahneman, D., 2011, Thinking, fast and slow, Farrar, Straus and Giroux

Kahneman, D., O. Sibony, and C. R. Sunstein, 2021, Noise. A flaw in human judgment, Little, Brown Spark

Lapuschkin, S., S. Wäldchen, A. Binder, G. Montavon, W. Samek, and K.-R. Müller, 2019, "Unmasking clever Hans predictors and assessing what machines really learn," Nature Communications 10(1096)

Pearl, J., 2000, Causality: models, reasoning and inference, Cambridge University Press

Pearl. J. and D. Mackenzie, 2018, The book of why: the new science of cause and effect, Effect. Basic Books.

Pumperla, M., and K. Ferguson, 2019, Deep learning and the game of Go, Manning

Puppe, F., 1993, Systematic introduction to expert systems, Springer Verlag

Rosenblatt, F., 1958, "The perceptron: a probabilistic model for information storage and organization in the brain," Psychological Review 65:6, 386-408

Roubini, N., 2022, Megathreats, John Murray Publishers

Rudin, C., 2019, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nature Machine Intelligence 1:5, 206-215

Samek, W., T. Wiegand, and K.-R. Müller, 2017, "Explaining artificial intelligence: understanding, visualizing and interpreting deep learning models," ITU Journal: ICT Discoveries 13 (Special issue no. 1)

Teoh, E. R., and D. G. Kidd, 2017, "Rage against the machine? Google's self-driving cars versus human drivers," Journal of Safety Research 63, 57-60

Tsaih, R.-H., H.-L. Chang, C.-C. Hsu, and D. C. Yen, 2023, "The AI Tech-Stack Model," Communications of the ACM 66:3, 69-77

Turing, A. M., 1950, "Computing machinery and intelligence," Mind LIX:236, 433-460

Wodecki, A. (2018). Value Creation and Competitive Advantage Models. In A. Wodecki, Artificial Intelligence in Value Creation (pp. 1-70). London, UK: Palgrave Macmillan

Zhao, W. X., et al. (2023). A Survey of Large Language Models, https://tinyurl.com/yjf68va5