Editorial

The golden age of Artificial Intelligence

Inaugural Editorial

Okyay Kaynak[1]

1 Introduction

During the last two decades, profound technological changes have occurred around us, supported by disruptive advances both on the software and hardware sides. Additionally, we have witnessed a cross-fertilization of concepts and an amalgamation of information, communication, and control technology-driven approaches. This has led to what is termed as digital transformation, i.e., the integration of digital technology into all areas of business, fundamentally changing how companies operate and deliver value to customers. The most recent development is integrating Artificial Intelligence (AI) in digital transformation as the primary enabler and the facilitator. It is expected that the applications of AI will truly transform our world and impact all facets of society, economy, living, working, healthcare and technology, creating a need for a journal that provides rapid dissemination and open access to research across diverse disciplines. This is the motivation behind the establishment of this new journal. I am honoured to be its Founding Editor in Chief.

2 The past

AI, as a term, is not new. It has been around us for more than 70 years. In Fig. 1, a timeline is shown with the critical events that have led to its present, glorious status, as well as the two AI "winters." In what follows, some of the bullets on the timeline will briefly be touched upon. Similar timelines with more details can be found in Ref. [1, 2].

2.1 Enigma

In 1941, Alan Turing, together with his colleague Gordon Welchman developed the Bombe Machine that deciphered the military codes used by Germany and its allies, popularly known as the ENIGMA code. The "enigma" was a type of enciphering machine. ENIAC (Electronic Numerical Integrator and Computer) was the first programmable, electronic, general-purpose digital computer developed at the United States Army's Ballistic Research Laboratory, mainly for calculating ballistic trajectories. It then cost more than US$ 500.000, in today's terms, close to 8 million dollars.

✉ Okyay Kaynak, okyay.kaynak@boun.edu.tr | [1]Bogazici University, Istanbul, Turkey.
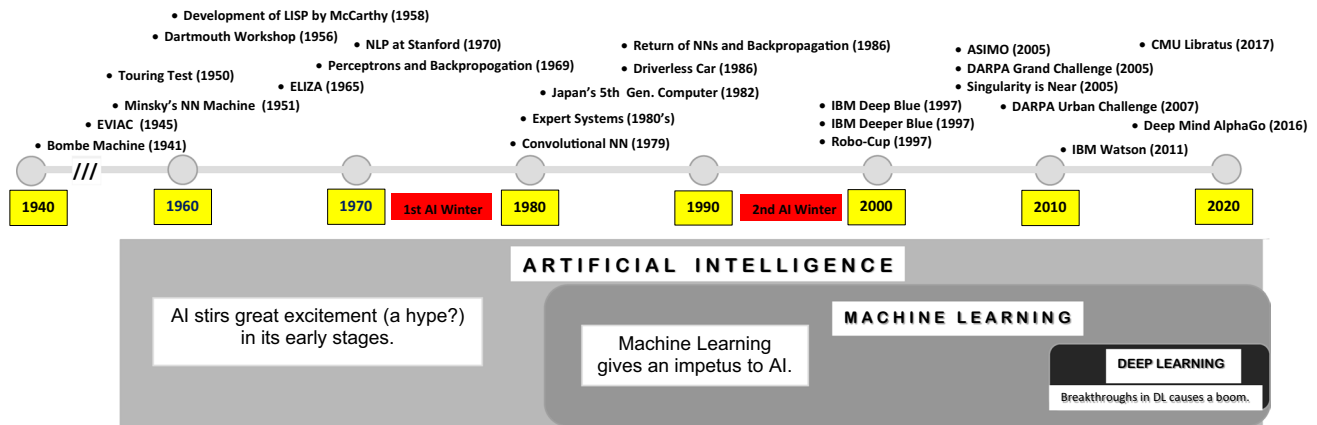
Fig.1 The timeline of AI

2.2 The turing test

Also known as the imitation game, is a test to determine whether a machine can exhibit intelligent behaviour. It has been discussed extensively in AI Community as a controversial topic, and it appears that it will continue to be controversial. Even today, there is no consensus on the answer to the question Turing posed; "can machines think [3]?" The reference here is to a video clip, illustrating in black and white, the state of AI in the 1960s, in which legendary AI pioneers like Jerome Wiesner, Oliver Selfridge, and Claude Shannon speak. It is amazing to listen to the hyped expectations that are, 60 years later, still to come!

2.3 Let there be AI! (Dartmouth Summer School 1956)

The term Artificial Intelligence was coined in 1955 by John McCarthy and his colleagues Marvin Minsky (Harvard University), Nathaniel Rochester (IBM), and Claude Shannon (Bell Telephone Laboratories) in the proposal for the funding of "2-month, a 10-man study of Artificial Intelligence." The workshop took place in the following year, in 1956, which is generally considered as the official birthdate of the new field. In their proposal, the proposers made very ambitious statements such as "(i) the study is to proceed based on the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it, (ii) an attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves."

2.4 LISP

In the 1960s, the most popular programming language used in AI research was LISP. McCarthy developed the basic ideas during 1956–1958, and it quickly became a common language for AI programming. The main reason for its popularity was that learning could be incorporated in LISP as self-modifying programs. The popularity of LISP was so high that special computers, so-called "LISP machines," that could run LISP programs efficiently and effectively swarmed the market. The most notable machine was PDP-10 of Digital Equipment Corporation (DEC). Even I used one in my early years of research!

2.5 The first AI winter

The expectations of the Dartmouth Summer School (as voiced in Ref. [3]) were far from being realistic and overly ambitious. The timing was just not right for many reasons; most notably, the necessary computing power was just not there. Millions of dollars were poured out to make the vision as described in the workshop proposal come true, but no tangible results! Consequently, in 1973, in response to the pressure from Congress, the U.S. and British

Governments stopped funding the undirected research into Artificial Intelligence. The following years would later be known as the "First AI Winter."

2.6 The second spring

Some years later, in 1982, an initiative by Japan's Ministry of International Trade and Industry (MITI), titled the Fifth Generation Computer Project (FGCP), reignited the AI torch. With eight top computer companies, billions of dollars were poured out to create a new generation of computers based on massively parallel computing and logic programming. The goal was to develop integrated systems, both in hardware and software, suitable for the computer applications in the shift from "information processing" to "knowledge processing." At about the same time, expert systems, composed of two subsystems, the inference engine and the knowledge base, became a commonly used term in the AI arena, used almost as a synonym for AI. They were the first truly successful forms of Artificial Intelligence software.

2.7 The autumn and the second AI winter

The goals of FGCP were too ambitious, again for the technical capacities and the capabilities of the computers of those years, and most of them could not be met. So much so that in the 1984 annual meeting of the Association for the Advancement of Artificial Intelligence (AAAI), Roger Schank and Marvin Minsky talked about the coming AI Winter. Soon after, the funding of the FGCP ceased, the final blow was struck by the collapse of Lisp Machines Inc (LMI), and AI fell out of the limelight.

3 The present

After having gone through two "winters" in the early seventies and the late eighties, AI made a glorious come back at the beginning of the second millennium, enabled by the emerging massive computing power, the collection of colossal data sets (the big data phenomenon) and the advances in data analytics (from descriptive to predictive, even to prescriptive). Investment and interest in AI boomed in the first decades of the twenty-first century when machine learning (ML) was successfully applied to many problems in academia and industry. The opportunities provided by deep learning (DL) intensified the boom.

The most notable cornerstones in the near past section of the timeline are the performances of the AI-enabled IBM machine Watson and the Google DeepMind machine AlphaGo. In 1997, another IBM machine Deep Blue had played against the chess champion Kasparov and drew. IBM formulated some improvements on its program, which was mainly expert system based (I call this machine "Deeper" Blue), and at another match, Kasparov lost. IBM Watson was designed more than a decade later as a question answering machine that incorporated more than 100 software technologies, mainly ML-based, running on an 80 teraflops supercomputer. Although Watson has been used for several purposes, it made the headlines when it defeated the US champions in the game Jeopardy [4]. However, a stupid mistake it made in learning from data was termed as artificial stupidity as an antonym to Artificial Intelligence!

As compared to chess, go is a much more complex board game; at the opening stage, there are 20 possible moves in chess, whereas in go, the number is 321. It, therefore, requires much more expertise and deep strategic thinking. AlphaGo is a software developed by Google DeepMind to play go and, in 2016, it defeated the legendary world champion Lee Sedol by four games to one.

In real world applications, there often are cases when an AI program must operate under conditions in which the information available is incomplete and/or the parties engaged may be hiding information or even engaging in deception. Libratus is an AI program developed at Carnegie Mellon University, designed to operate under such conditions. It proved its capabilities in 2017 by winning against four top-class human poker (no-limit Texas hold 'em) players in a 20-day competition. The developers intend to use Libratus in applications that involve complex decision making based on imperfect information such as setting military strategies, negotiating business deals, or planning a course of medical treatment.

4 The future

What has been narrated above indicates that the developments in AI have been happening at an increasing rate of change. Presently we are at the stage of Artificial Narrow Intelligence (ANI), but by the end of the present decade, we may be more into Artificial General Intelligence (AGI). An exciting prospect! The possible path to the future will be discussed in a full-length article titled "Qua Vadis AI?" authored by my colleagues and myself. It is planned to appear in this journal very soon; just follow us!

The industrial applications of AI have not been discussed much in literature (except perhaps in the recent EU white paper as will be discussed later) but very much deserves attention. In modern large-scale industrial processes, in parallel with the move to Circular Economy from a linear one, there are increased demands for fuel efficiency, conservation of resources, cost and energy savings, and other similar requirements for plant-wide optimization. Conventional methods of plant optimization, fault diagnosis and control require system models, which are established either analytically starting from first principles or through detailed identification techniques. The feasibility and complexity of these approaches vary significantly among specific applications, and the monitoring and control performances rely heavily on the precision of the model. Furthermore, the model established is likely to change in time, with changing operating conditions.

The power of data-enabled or, more generally, data-driven decision-making and analytics have already helped us achieve enhanced control and operations for large-scale industrial processes. In the last 25 years, multivariate statistical process data analytic tools have been effectively used in process operations and control. In recent years, however, increasing attention is seen in the process industry in AI-augmented data-driven techniques. The intent is to efficiently extract the information available in the process and quality data to achieve three major tasks, process data analytics, prediction, and prescriptive actions.

5 The Journal

As has already been stated, AI has become very popular in all parts of the world and in both academia and industry and governments. It is projected to be a core skill for the future, and its capabilities are expected to be utilized in almost everything (the AI market is projected to grow to $190 billion by 2025). "Adopt AI" is nowadays a very commonly used slogan. Not a single day passes without the announcement of a new and exciting AI application. As I am writing this editorial, I have done some parallel processing and read a recent article on "caregivers" [5]!

It is stated in the recently published EU white paper [6] that "Artificial Intelligence is developing fast. It will change our lives by improving healthcare (e.g., making diagnosis more precise, enabling better prevention of diseases), increasing the efficiency of farming, contributing to climate change mitigation and adaptation, improving the efficiency of production systems through predictive maintenance, increasing the security of Europeans, and in many other ways that we can only begin to imagine. At the same time, Artificial Intelligence (AI) entails a number of potential risks, such as

Fig. 2 Number of publications per year
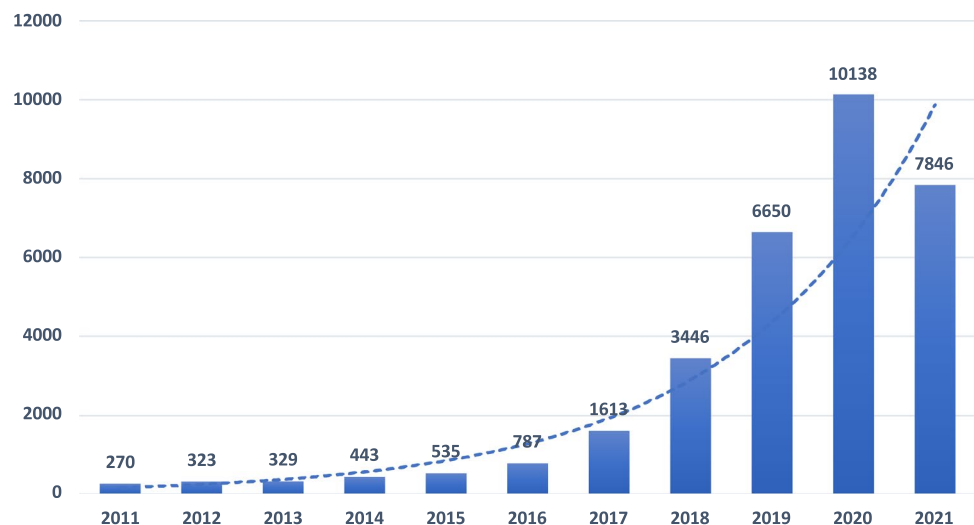
Fig. 3 Number of publications (top 10 countries)
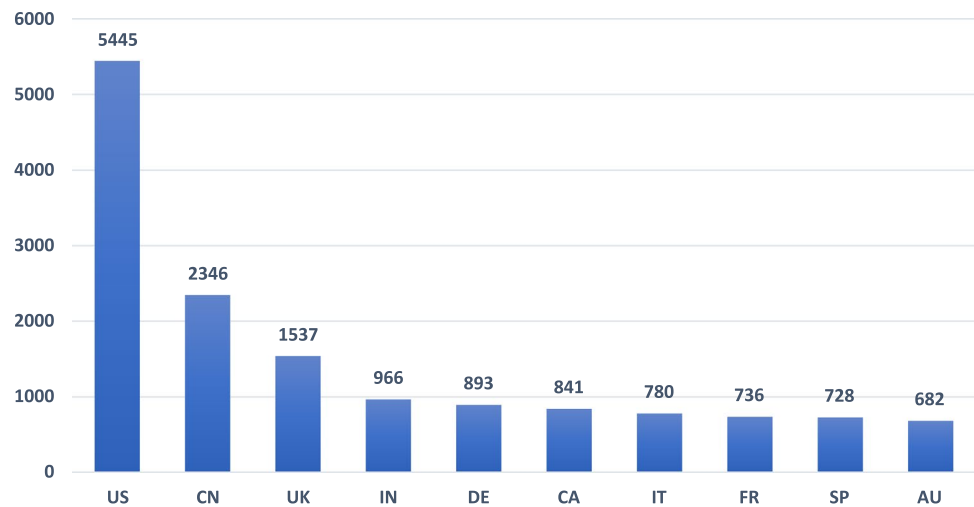
opaque decision-making, gender-based or other kinds of discrimination, intrusion in our private lives or being used for criminal purposes."

It is, therefore, no wonder that there has been an explosion in the number of articles published in literature. The results of a search in the Web of Knowledge (all databases) with the search string "Artificial Intelligence" are depicted in Figs. 2 and 3. Such a surge in the number of publications translates into 5–10 times more submissions to the existing journals that inevitably results in some bottlenecks on the journey from submission to publication, forcing researchers to seek new avenues for dissemination of their work. Discover Journals series by Springer, which is a collection of fully open access journals, are committed to providing all authors a streamlined submission process, rapid review and publication, and a high level of author service at every stage. The related link is https://www.springer.com/gp/campaign/discover-journals. Discover Artificial Intelligence is one of the journals in the series that covers both the theory and the applications of AI in areas such as industry, healthcare and medical diagnostics, transport, agriculture, education, and economics. It also covers AI as it relates to machine learning and deep learning, data analytics, knowledge reasoning and discovery, natural language processing, computer vision, robotics, as well as social sciences, ethics, legal issues, and regulation. Its editorial board is composed of very prestigious names in three tiers; I am honored the hold the Editor in Chief position; in the tier below are the Deputy Editor in Chiefs and in the third tier are our Associate Editors, not to mention the managerial staff (the Managing Editor Yina Liu and her team) that are committed to ensuring the smooth and timely operation of the journal by providing a high level of author service at every stage. Please visit the journal website for more information on the scope and the editorial board https://www.springer.com/journal/44163.

6 The contents of the inaugural issue

The inaugural issue is composed of seven papers, three were submitted in response to our publicity campaign, and four were submitted by the members of the Editorial Board. Regarding their types, they are diversified. Two are perspective papers, two are reviews, two report research results, and one is a case study. A brief description of them, limited to a sentence or two, is as follows:

The opening article of the issue presents a perspective, titled "Ethical and Legal Responsibility for Artificial Intelligence", in which the author Patrick Henz, discusses the "nature" of Artificial Intelligence, including the risks its posing, and who is responsible for systematic errors, from an ethical, but also legal point [7].

The next article by Deborah Petrat, titled "Artificial Intelligence in Human Factors and Ergonomics—An overview of the current state of research" reviews the topics of human factors and ergonomics, so that a smooth implementation of AI applications can be realized. To map the current state of research in the area, three systematic literature reviews with different focuses are conducted [8].

A Generative Adversarial Network (GAN) is a class of machine learning framework in which two neural networks contest with each other. GANs have become very popular in recent years, especially in image analysis, as they open some exciting new ways for medical image generation, expanding the number of medical images available for deep

learning methods. In the article titled, "When Medical Images Meet Generative Adversarial Network: recent Development and Research Opportunities," the authors discuss the topic in depth [9].

The fourth article of the issue is titled "Diabetes and Conversational Agents: the AIDA Project Case," in which the authors introduce their Artificial Intelligence Diabetes Assistant (AIDA). It consists of a text-based chatbot and a speech-based dialog system [10].

In the next article, the authors Ambareesh Ravi, Fakhri Karray present a perspective, titled "Exploring Convolutional Recurrent Architectures for Anomaly Detection in Videos: a comparative study," the authors explore a variety of Convolutional Recurrent architectures and the influence of hyper-parameters on their performance for the task of anomaly detection [11].

The following article is related to Industrial Cyber Physical Systems (ICPS) in which large volumes of data are generated at high velocity. The generated data streams are susceptible to dynamic and abrupt changes which are formally defined as concept drifts. In the article, titled, "Continuous Detection of Concept Drift in Industrial Cyber-Physical Systems using Closed Loop Incremental Machine Learning" the authors propose an unsupervised, self-adaptive machine learning algorithm for continuous concept drift detection in industrial CPS [12].

Applying machine learning methods to improve the efficiency of complex manufacturing processes, such as material handling, can be challenging. The last article of the issue, titled "Applying Reinforcement Learning to Plan Manufacturing Material Handling," addresses this issue. The authors demonstrate the applicability of reinforcement learning with a multi-objective reward function to realistically complex material handling processes [13].

7 The final words (invitation to contribute)

I am aware that this editorial for the inaugural issue has perhaps become too long. I hope you have found it worth reading. The video on "The Thinking Machine" [3] should perhaps help us to keep a clear head in all the AI hype surrounding us. I am also aware that the web pages related to the journal are full of information. If you get lost anywhere, please do not hesitate to contact me or the Managing Editor, Yina Liu. We would like to make our difference by at least being easily reachable.

I would like to conclude by inviting you to contribute to our journal, especially in the form of topical collections, which are guest-edited special issues, on emerging hot topics of relevance to all aspects of Artificial Intelligence. The details on how to propose one are available at the link, https://www.springer.com/journal/44163/updates/18964808. Please contact us if you have one in mind. Furthermore, we welcome your research outputs that present new ways of looking at AI problems and applications and demonstrates the value and the effectiveness of their implementation in modern society.

Okyay Kaynak
Editor-in-Chief

References

1. Gil Press G. A Very Short History Of Artificial Intelligence. Forbes. https://www.forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/?sh=2d73376a6fba. Accessed 9 Sept 2021.
2. Anyoha R. A brief history of Artificial Intelligence: on the past, present, and future of Artificial Intelligence". Calif Manage Rev. 2019;61(4):5–14.
3. The Thinking Machine (Artificial Intelligence in the 1960s). https://youtu.be/aygSMgK3BEM. Accessed 9 Sept 2021.
4. Markov J. Computer wins on "Jeopardy!": trivial, it's not. New York: The New York Times; 2021.
5. Campbell M, Jovanovic M. Conversational Artificial Intelligence: changing tomorrow's health care today. Computer. 2021;54(8):89–93.
6. Commission E. White paper on Artificial Intelligence—a European approach to excellence and trust. Brussels: European Commission; 2020.

7. Henz P. Ethical and legal responsibility for Artificial Intelligence. Discov Artif Intell. 2021. https://doi.org/10.1007/s44163-021-00002-4.
8. Petrat D. Artificial Intelligence in human factors and ergonomics—an overview of the current state of research. Discov Artif Intell. 2021. https://doi.org/10.1007/s44163-021-00001-5.
9. Li X, Jiang Y, Rodriguez-Andina JJ, Luo H, Yin S, Kaynak O. When medical images meet generative adversarial network: recent development and research opportunities. Discov Artif Intell. 2021. https://doi.org/10.1007/s44163-021-00006-0.
10. Alloatti F, Bosca A, Caro LD, Fabrizio F. Diabetes and conversational agents: the AIDA project case study. Discov Artif Intell. 2021. https://doi.org/10.1007/s44163-021-00005-1.
11. Ravi A, Karray F. Exploring convolutional recurrent architectures for anomaly detection in videos: a comparative study. Discov Artif Intell. 2021. https://doi.org/10.1007/s44163-021-00004-2.
12. Jayaratne D, Silva DD, Alahakoon D, Yu X. Continuous detection of concept drift in industrial cyber-physical systems using closed loop incremental machine learning. Discov Artif Intell. 2021. https://doi.org/10.1007/s44163-021-00007-z.
13. Govindaiah S, Petty MD. Applying reinforcement learning to plan manufacturing material handling. Discov Artif Intell. 2021. https://doi.org/10.1007/s44163-021-00003-3.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.