

IS THIS ARTIFICIAL INTELLIGENCE?

Vladan Devedžić

University of Belgrade, Faculty of Organizational Sciences, Belgrade, Serbia

Abstract. *Artificial Intelligence (AI) has become one of the most frequently used terms in the technical jargon (and often in not-so-technical jargon). Recent advancements in the field of AI have certainly contributed to the AI hype, and so have numerous applications and results of using AI technology in practice. Still, just like with any other hype, the AI hype has its controversies. This paper critically examines developments in the field of AI from multiple perspectives – research, technological, social and pragmatic. Part of the controversies of the AI hype stem from the fact that people use the term AI differently, often without a deep understanding of the wider context in which AI as a field has been developing since its inception in Mid 1950s.*

Key words: *Intelligence, Artificial Intelligence (AI), technology, applications, reality check.*

1. INTRODUCTION

Artificial Intelligence (AI) is seeing an unprecedented rise in popularity for more than a decade. Several traditional subfields of AI have developed almost to the level of disciplines per se, and there are more and more practical applications of different technologies that have been developing for years under the AI umbrella. This has affected many sectors, and has attracted attention of not only technology developers, but also of educators, social scientists, artists, governments, media and wider public.

On the other hand, there are many apparently simple questions that are still waiting for appropriate answers. *What exactly is AI, in the first place? How intelligent is an intelligent system? What are the criteria to call a system an AI system, or an intelligent system?*

In order to set the stage for discussing these questions further, a brief review of some real-world examples of systems and applications called AI is a good starting point.

Spam filtering is one of the commonly known examples of applying AI in email services, but it's less commonly known that smart email categorization and labelling is also AI-powered [1]. Even fewer email users are aware of AI behind smart replies, nudging which emails they haven't answered or ignored.

Received August 20, 2020

Corresponding author: Vladan Devedžić

University of Belgrade, Faculty of Organizational Sciences, Jove Ilića 154, 11000 Belgrade, Serbia

E-mail: devedzic@gmail.com

AI voice-to-text apps for smartphones, like Speechnotes¹ and Voice Notebook², can convert speech to text and can also convert an audio file to text. The same technology powers smart personal assistants, like Google Assistant³, Alexa⁴ and Cortana⁵, that can perform Internet searches, set reminders, integrate with your calendar, create to-do lists, order items online and answer questions (via Internet searches).

When Google Maps recommends the fastest route through a city on someone's smartphone, it intelligently takes into account not only the traffic speed, but also the road construction, accidents and different user-reported conditions [2]. Likewise, ride-hailing-and-sharing apps like Uber can accurately calculate the price of a ride, predict the passenger's demands, determine optimal pick-up locations and even compute the estimated time for food delivery [3]. Some will be surprised to learn that AI autopilots on commercial flights are in charge of flying the aircraft for most of the flight time – human-steered times are typically just during takeoff and landing [4].

And that easily shifts attention to self-driving cars, buses and trucks, a largely debated AI topic that until very recently referred only to experimentation that used to spark our imagination, but nowadays is slowly becoming a reality [5]. These vehicles are smart enough to drive at an optimal speed, to follow the signs, to pay attention to the stop lights, pedestrians and other cars, and safely bring the passengers and loads to their destinations.

Using AI-enabled technology in military applications has always been one of the driving forces in developing AI further. Typical current applications include unmanned (self-driving) vehicles, combat robots, drone swarms and autonomous action [6], [7]. They allow for running dangerous, suicidal missions, and have opened a whole new line of military strategy and tactics development. A good recent example that uses the AI techniques, called adversarial machine learning [8], is model turtle created at MIT – a robot that looks like a turtle to humans, but can easily fool other AI-powered robots and surveillance drones, to which it looks like a rifle [9]. This leads to a series of adversarial algorithmic camouflage tactics, like hiding military planes, tanks, and other objects, “blinding” missiles, and so on.

Image recognition and face recognition systems have become quite popular. Facebook⁶ highlights faces on an uploaded image and suggests the user friends to tag, using AI to recognize faces. Snapchat⁷ goes into a slightly different direction – it can also track facial movements. Similarly, Instagram⁸ uses AI to identify the contextual meaning of emoji. Amazon Rekognition⁹ can recognize faces of celebrities, and so can Microsoft Azure Custom Vision¹⁰ image recognition cognitive service. Google Cloud Vision¹¹ and Amazon Rekognition are currently among the leaders of general object recognition and content

¹ <https://speechnotes.co/>

² <https://voicenotebook.com/>

³ <https://assistant.google.com/>

⁴ <http://alexa.amazon.com/spa/index.html>

⁵ <https://support.microsoft.com/en-us/help/17214/windows-10-what-is>

⁶ <https://www.facebook.com/>

⁷ <https://www.snapchat.com/>

⁸ <https://www.instagram.com/>

⁹ <https://aws.amazon.com/rekognition/>

¹⁰ <https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/>

¹¹ <https://cloud.google.com/vision/>

detection on images. Google Lens¹² brings up relevant information related to objects it identifies using visual analysis, Fig. 1.



Fig. 1 The photo of the author's desk taken by the Google Lens app run by his smartphone (left) and part of the information shown by the app (correctly except for the color) as a result of the AI-based image analysis (right)

In the banking sector, fraud detection platforms based on machine learning (ML), such as the one created by the Teradata¹³ firm, are in high demand [10]. They are capable of recognizing potential fraud transactions by differentiating between acceptable deviations from the norm and critical ones. Acceptable deviations are treated as false positives, so the system can “learn” from its mistakes. The data used to train the ML model include recent frequency of transactions, transaction size, geolocational data, the kind of retailer involved, etc.

So, what is it in these (and many, many more) systems and applications that is most often called AI?

2. DEFINING AI?

The question mark in the subheading is intentional. AI is notoriously hard to define – in fact, there are many definitions and none of them is dominant in the AI community; P. Marsden has compiled a list of a few dozens of popular definitions [11]. Extracting and mixing bits and pieces from several of them, in this article AI is understood primarily as technology capable of exhibiting skills typically associated with human intelligence, such as the ability to perceive, learn, reason, abstract (classify, conceptualize and generate rules) and act autonomously. It is also the science and engineering of creating such technology, where intelligence is the computational part of it that enables machines to exhibit behaviors and actions that would be called intelligent if a human were so behaving, i.e. that would require intelligence if they were done by humans.

An important characteristic of an AI system is that it can figure out things for itself, and then act based on that information. The most popular textbook on AI [12] stresses a

¹² <https://lens.google.com/>

¹³ <https://www.teradata.com/>

variation of that characteristic: “AI is the study of agents that receive percepts from the environment and perform actions... a rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best-expected outcome,” i.e. has the ability to achieve goals in the world in an optimal way.

There are at least two distinct points in this understanding/description: (a) AI is technology, more precisely computational technology; and (b) it behaves and acts in a way that is typically associated with human intelligence. What makes things slip away in all attempts to define AI is not part (a); it is part (b).

2.1. What is *intelligence*?

The much-quoted line of R.J. Sternberg that “viewed narrowly, there seem to be almost as many definitions of intelligence as there were experts asked to define it” [13] reveals in a concise way that all attempts to define intelligence are inherently controversial. And, just like in the case of defining AI, there are collections of definitions (e.g., [14]) and broad statements and commentaries that outline only vague conclusions about the nature of intelligence, its origins and current scientific evidence.

This article adopts two broad statements of this kind, which describe intelligence as:

“A very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience. It is not merely book learning, a narrow academic skill, or test-taking smarts. Rather, it reflects a broader and deeper capability for comprehending our surroundings – “catching on,” “making sense” of things, or “figuring out” what to do.” [15]

“Ability to understand complex ideas, to adapt effectively to the environment, to learn from experience, to engage in various forms of reasoning, to overcome obstacles by taking thought... Concepts of “intelligence” are attempts to clarify and organize this complex set of phenomena.” [16]

Note, however, that all such statements and attempts to define (or, at least, characterize) intelligence can lead to a vicious circle. One now needs to define each of these abilities, like understanding, thinking, reasoning, learning, adapting, etc. This is equally difficult as defining intelligence, since “although considerable clarity has been achieved in some areas, no such conceptualization has yet answered all the important questions, and none commands universal assent” [16]. Moreover, there can be substantial individual differences in performance related to these complex abilities, and they can vary even for the same person in different domains, under different circumstances, and so on. Mechanisms to measure this performance do exist (e.g., IQ), but judgement can be based on different criteria.

2.2. How intelligent is an AI system?

Given the extremely high complexity of intelligence itself and of the abilities associated with it, developers of AI systems typically focus only on some narrow aspects of intelligence or to a specific dimension of intelligence, such as knowledge representation, reasoning, learning, and image analysis and interpretation.

Unfortunately, this can lead to big differences in judging how intelligent is an AI system.

2.2.1. The Turing test

As early as 1950, Alan Turing suggested that a program/machine should pass a behavioral intelligence test if it was to be called intelligent [17]: it should have a 5-minute typed-messages conversation with a human interrogator, and the interrogator then has to guess if the conversation was with a program or with a person; the program/machine passes the test if for at least 30% of the time the interrogator believes she/he is making this conversation with a person [12].

The modern-time interpretation of the Turing test [12] is that such a program/machine should be able to communicate successfully with the interrogator using a natural language, should be capable of representing and storing information and knowledge about what it hears and using that knowledge for reasoning when answering questions and drawing conclusions. In addition, it should be able to learn new knowledge and patterns and to adapt to new situations, as well as to perceive objects using its sensory input and manipulate the objects accordingly (robotics).

Ever since the Turing test was proposed, it has created intense debates. Philosophers have argued that there are things that machines cannot do, others have cited mathematical proofs that some questions are in principle unanswerable by formal systems, and some strongly support the stance that human intelligence is much too complex to be captured by machines. However, in recent years there have been several announcements about AI systems passing the Turing test [18], [19], [20]. These typically initiate counter-arguments and stay confined to academic circles; so far, there has been no much reaction from technologists.

2.2.2. Weak AI vs. strong AI

Weak AI systems are those that can act *as if* they were intelligent, i.e. they can *simulate* human cognitive function. They can only *appear* to think, but definitely lack consciousness. They can follow certain rules and pre-programmed behaviors – even complex ones – but cannot do anything beyond these rules and behaviors. For example, a chess-playing program cannot be used as a personal assistant and vice versa. As J. Searle puts it [21]: “According to weak AI, the principal value of the computer in the study of the mind is that it gives us a very powerful tool. For example, it enables us to formulate and test hypotheses in a more rigorous and precise fashion. But according to strong AI, the computer is not merely a tool in the study of the mind.”

In contrast to weak AI, the hypothesis of strong AI is that an AI system should actually *have* human cognitive abilities and states, not just simulate them. Strong AI is not about building tools that help test psychological explanations; it is about building systems that “are themselves the explanations” [21]. In other words, according to strong AI, intelligent programs should have their own autonomous perception, beliefs, emotions, and intentions; they should *be* minds.

Current systems called “AI systems” are typically developed with the weak AI hypotheses in mind [12]. Developers are happy if their programs work, and do not care much if people call them real intelligence or just simulated one.

A related problem is the level of sophistication of an AI system. Current AI systems can easily beat even the best human players in computer games or in the games of chess and Go, but can neither understand nor feel the meaning of fairy tales and stories for young children [22], let alone capture their bottom-lines and morals.

2.2.3. AI effect

Critics of weak AI often discount a successful AI technology with not viewing it as being real intelligence, regardless of the fact that it was once considered AI [23]. This is called AI effect: before the technology becomes part of everyday life, i.e. before it comes out from the confines of AI research labs, it has a special aura; it looks magic and truly intelligent. Once it is better understood by the majority and gets built into products and tools used by many, the thrill is gone – it loses the ‘AI’ label and becomes just technology.

As a side effect, advancements in technologies that have once lived under the AI umbrella sometimes make these technologies break away from the ‘AI’ label and get re-branded: expert systems have come out of the AI auspices and become a technology per se, artificial neural networks are often called just neural networks, and everybody says just chatbots, not AI chatbots.

Some see the cause of AI effect in the difference between the strong AI and weak AI concepts [24]. Those who are ready to remove the ‘AI’ label from technology originating from AI research typically align themselves with the strong AI approach: if an AI problem has been solved, it’s no longer AI; true AI is a problem that has not been solved yet. A possible way out is to take a different perspective: since AI today is typically weak AI, perhaps a down-to-Earth question to ask is “Can a specific problem be solved with weak AI or not?”

It is also a good idea to occasionally “see the world differently” – what do researchers in other, more-or-less related disciplines, have to say about intelligence?

2.3. Intelligence seen from different research perspectives

There is a dichotomy in explaining AI from technological and other perspectives. While technology-centered AI development focuses on systems that work accurately and fast, have exciting functionality and demonstrate certain aspects of intelligent behavior, experts in other disciplines are more interested in advancing the understanding of the phenomenon of intelligence.

2.3.1. Neuroscience

Neuroscientists have made some progress identifying various neurological factors relevant for intelligence [25]. It is now known that intelligence and functioning of the brain are related to the overall brain volume, cortical thickness, white matter volume, grey matter volume, white matter integrity, neural efficiency, etc. But it is also known that such factors are only partly responsible for differences in intelligence among different humans (as well as among different members of other species). Popular techniques/technologies used in non-invasive scanning of human brain include electroencephalography (EEG), magnetic resonance imaging (MRI), functional MRI (fMRI), etc. For example, recent uses of powerful MRI scanners have enabled analysis of functional units inside the layers of the human cortex (responsible for high level of cognition) and seeing for the first time how information flows between collections of neurons in a live human brain [26]. Note that this is extremely important for neural network research in AI – neural networks as we know them today are *models* based on never-proved *assumptions* of how neurons exchange information. Moreover, such scanners have brought neuroscientists one step

closer to understanding of human memory. Likewise, an analysis of over 18.000 MRI scans of people over 44, paired with four cognitive tests from the UK Biobank study has revealed that the brain size has only a minor correlation with intelligence, biological sex has no impact on intelligence, and intelligence is largely influenced by different brain regions [27].

Note, however, that neuroscientists admit that although now we have considerably more evidence about how human brain functions and what regions of the brain are responsible for intelligence, we still don't know what intelligence really *is*; a lot of further research is needed to fully understand it. That's why some neuroscientists take a different approach. Due to the fact that human brain is extremely complex, they make attempts to understand how the brain of simpler species works. For example, a notable success has been achieved with studying the brain of fruit flies (*Drosophila melanogaster*) using electron microscopy – the entire brain of an adult female fly has been imaged at synaptic resolution [28]. However, a fact very relevant for AI research is that in spite of now having an unbiased mapping of synaptic connectivity of the fruit fly, *synthesizing* its brain – the size of a poppy seed – is not even at sight.

2.3.2. Psychology

Research and experiments in cognitive psychology have led to theories about how humans represent knowledge and how they process it in order to make inferences and decisions, create explanations, analyze situations at hand, reach conclusions and so on. The knowledge represented pertains both to external world and to internal mental states, like beliefs, emotions, attitudes and desires [29]. Information perceived from the world (both external and internal) gets encoded into mental representations and is either processed immediately, or is stored in memory for later retrieval and processing.

There are several basic forms of mental representations: spatial (e.g., the placement of objects in a room), feature (such as dogs *bark*, *can run*, *have four legs*, *are faithful*,...), network (like *Irish setter is a setter*, *Irish setter is red*, *Irish setter has bird sense*, *Irish setter is a dog*, *dog is an animal*), and structured (like *a plate is on the table*, *a drawer is under the table*, *the drawer is closed*,...). These forms themselves have their structures. There are also specific processes associated with each form, capable of accessing and using information and knowledge represented within a specific form. For example, in the network representation example shown above, the *is a* relation between *Irish setter* and *dog* enables accessing *dog* features indirectly and inferring that *Irish setter* can *bark*.

A powerful tool of human thought processes is *abstraction*. It enables ignoring some information (i.e., not representing it, abstracting it away). This is very important in terms of the efficiency of processing the information that *did* get stored within the representation – it can be found and accessed more quickly, since the search space is more compact without the information that got abstracted away.

Cognitive science lays the bridge between cognitive psychology and AI. It develops computational models of different forms of mental representations and their related processes. Note, however, that these models only *theoretically mimic* human thought. In reality, we know very little about how knowledge is represented and processed in human brain [30], in spite of valuable recent discoveries like the one that has revealed the brain's code for facial identity [31]. Researchers are only beginning to tackle important problems

like the relation between consciousness and intelligence [32] and the one between intentionality and intelligence [33].

2.3.3. Philosophy

Ever since the inception of AI, philosophers have been intrigued with it. The already mentioned work related to strong AI ([17], [21], [22]) is but a tiny bit of discussions on the topic. Chapter 26 of [12] surveys philosophical pros and cons related to AI in much more details.

Some of the more recent considerations and debates in this area include V. Vinge's notion of (*technological*) *singularity* [34], built upon the earlier I.J. Good's concept of *intelligence explosion* [35]. Essentially, singularity means that if humans can create intelligence smarter than their own, then it could do the same, only faster. The concept has been further explored by R. Kurzweil [36], who projected that, given the pace of technological development, by Mid 2040s global computing capacity will exceed the capacity of all human brains, which will be a precondition for singularity.

Numerous philosophical speculations and debates have followed, on the grounds that human brains cannot even comprehend such a superior intelligence. Some have expressed fear that singularity can ultimately lead to the extinction of humans. Others strongly oppose this view, arguing that humanity has already entered "a major evolutionary transition that merges technology, biology, and society, where digital technology has pervaded the fabric of human society to life-sustaining dependence", transition that will ultimately lead to *Real AI (RAI)*, as "a globally distributed hybrid cyber-physical human intelligence converging all the emerging technologies: RAI = World Big Data + AI + ML (DNNs) + Cloud AI + Edge AI + IoT + 5G + Blockchain + Autonomous Things + Self-Driving Cars + Virtual Reality and Augmented Reality + 3D Printer + Quantum Computing + Smart Spaces + ..." [37]. Notably, natural intelligence is included in the concept of RAI.

Yet other opinions exist, expressing the view that intelligence might be simpler than we think [38], since the way that humans perceive the world is hierarchical in nature, relying on simple patterns at the lower levels and increasing in complexity at the higher ones [39]. This is to say that the essence of perception, thinking, reasoning and other intelligent processing is actually pattern recognition – a long studied area in AI. All RAI is viewed as a combination of a) relations/patterns/causality between entities in the environment, b) representation of a), and c) perception, cognition and reasoning in order to establish understanding of the environment and provide rational interaction with it. To this end, P. Domingos has introduced the concept of *master algorithm* [40], as a blend of different approaches to strong AI and to ML in particular – symbolic, connectionist, evolutionary, Bayesian and analogy-based – where different ML algorithms synergistically contribute to an asymptotically perfect understanding the world, the brain and intelligence.

Philosophers also study higher-level concepts and their relations to intelligence, starting from the much quoted and thought-inspiring book *Gödel, Escher, Bach: An Eternal Golden Braid* by D. Hofstadter [41]. These include deep links between art, music, creativity, algorithms, imagination and abstract math, subtly reflected in and subsumed by intelligence. For example, S. Mahadevan has proposed the new concept of *imagination machines* as "a powerful launching pad for transforming AI" beyond the "current realm of learning probability distributions from samples" [42]. Using numerous examples from

arts, literature, poetry, and science, he envisions a new field of study in AI, *imagination science*, where researchers would explore various ways of automating tasks like “generating samples from a novel probability distribution different from the one given during training; causal reasoning to uncover interpretable explanations; or analogical reasoning to generalize to novel situations”.

3. CURRENT FOCUS IN AI

Given the difficulties in setting the scope and the boundaries of AI, reconciling somewhat different approaches to it when it’s seen from the perspective of scholars of different backgrounds, as well as in resolving controversies that surround it, a pragmatic approach is to focus on its most popular subareas (at any given point in time).

At the time of writing this article (July-August 2020), the “popularity bar graph” of these subareas, published at the AI Topics¹⁴ Website (curated by the highly authoritative Association for the Advancement of Artificial Intelligence, AAAI¹⁵), looks as in Fig. 2. The popularity is measured by the number of entries in the AI Topics repository, related to specific topics.

It is obvious that ML is currently the most popular subarea of AI – out of the total of 336.000+ entries, about 160.000 are tagged ML. There are two major reasons for that. One of them is the flood of data that applications, businesses, different institutions, social networks, etc. generate. People want to make sense out of this extremely vast amount of data in order to improve their businesses and other activities, and ML comes as a rescue – it enables building a mathematical model based on sample data, known as “training data”, in order to make predictions or decisions with previously unseen data, but without being explicitly programmed to do so [43]. To build models and make predictions, ML closely relies on computational statistics, mathematical optimization and exploratory data analysis; thus, it is also referred to as predictive analytics. The models themselves come in various forms, such as neural networks, regression analysis, decision trees, support vector machines, etc.

Drilling down the graph shown in Fig. 2 reveals that out of the nearly 160.000 ML entries about 54.000 are related to neural networks (NNs), and about 32.000 are related to statistical learning. Among the different types of neural networks, currently most popular ones are *deep neural networks* (DNNs) that enable so-called *deep learning* (DL) [44], [45], [46]. Important types of DNNs include: *convolutional neural networks* (CNNs, typically used for image analysis, facial recognition, visual search, etc.) [44], [45]; *recurrent neural networks* (RNNs, useful in natural language processing, speech analysis, text analysis and so on) [44], [45]; and *generative adversarial networks* (GANs, often used to generate examples for image datasets, photographs of human faces, realistic photographs, cartoon characters and face frontal views, as well as to perform image-to-image translation, text-to-image translation, semantic-image-to-photo translation, and more) [47].

¹⁴ <https://aitopics.org/>

¹⁵ <https://aaai.org/>

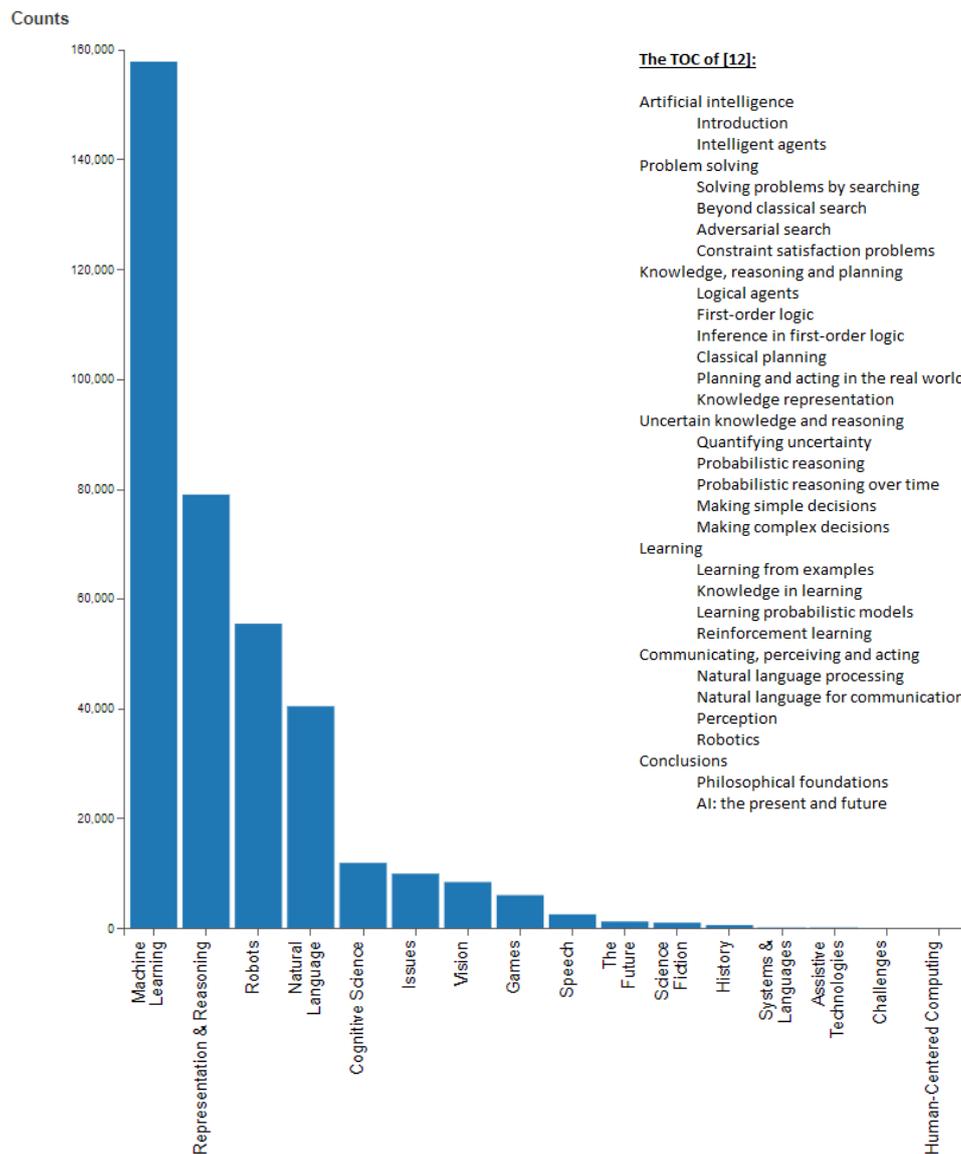


Fig. 2 The bar graph of popular AI topics at the time of writing the article (source: AI Topics Website, <https://shorturl.at/sAU28>) and the parts/chapters of the most popular AI textbook [12] (right)

The other reason for ML being so popular nowadays is the computational power of current ML technologies. The idea of learning new knowledge from data has been attractive in AI for decades, but only recently the computing technology has advanced to the level that has made it at least partially possible. Where it is not easily possible – e.g.,

requires too long processing time to build models that make predictions with a satisfactory level of accuracy – special-purpose computer hardware is usually the best solution. It can be a costly one, but it's a situation that further accelerates hardware development.

It should be also noted that ML and especially DNNs have become pervasive in other popular subareas of AI indicated in Fig. 2, notably in natural language processing (NLP) and in robotics. In NLP, application of DNNs has led to many advancements in language modeling, capturing semantic properties of words, natural language generation, machine translation, word- and sentence-level classification, sentiment analysis, and more [48]. In robotics, detection and perception of objects, robotic grippers, fine grasping and object manipulation, scene understanding and sensor fusion, as well as collision avoidance, are all greatly improved with careful use of DNNs [49].

The bar graph shown in Fig. 2 is actually much more accurate than the current, informally established public view of AI. This public view can be often seen in media and in popular press, blog posts and forums all over the Web: $AI \equiv ML$! A very frequent modality is AI/ML, and so is a less inaccurate “AI and ML”. There are also variations in a bit narrower scope, like ML/NN, ML/DL and the like. This has prompted more knowledgeable people to spawn all over the Web a series of images like the one on the left in Fig. 3, depicting the subsumption relationship between AI, ML and DL. However, the diagram on the right in Fig. 3 captures more details from the above discussion.

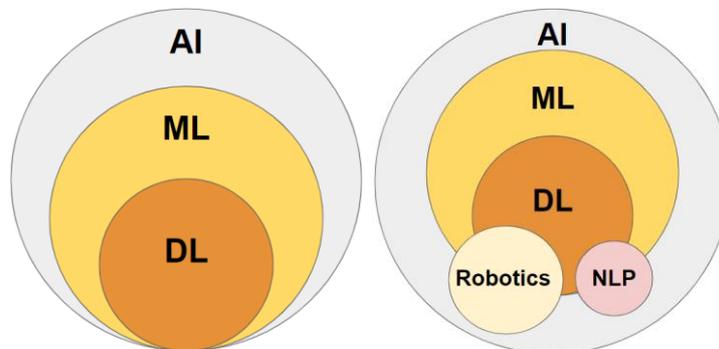


Fig. 3 Relationship between AI, ML and DL (left; after [50]) and a more detailed view based on the bar graph from Fig. 2 (right)

The righthand side of Fig. 2 shows the table of contents of the most popular AI textbook today, *Artificial Intelligence – A Modern Approach* [12]. Note that there is only a minor overlap with the bar graph on the left side. It further explains the diagram on the right side of Fig. 3 – many of the remaining topics still *are* part of AI (the outer circle in Fig. 3), but they are not in focus (which usually means lack of funding as well). A notable exception to this end is the broad subarea of AI – representation and reasoning (the second highest bin of the bar graph in Fig. 2). It has always been, and still is, in the core of AI. AI textbooks typically discuss only classical topics from this subarea (propositional logic, predicate logic, production rules, reasoning with uncertainty, fuzzy logic and systems, probabilistic reasoning and the like). However, there is a thriving research there as well (although it still

does not manage to catch much of the public attention) – new representation techniques and new efficient reasoning mechanisms have been devised recently [51], [52]. These largely pertain to topic modeling, knowledge graphs, conceptual modeling, representation of different types of thinking, knowledge interwoven with imperfect data, semantic summarization, the tradeoff between expressiveness and tractability, and constructing explanations.

The AI Topics Website largely reflects the views and interests of the AI community. However, views from other communities also matter. For example, Fig. 4 shows economic perspective on strategic development of AI. ML is still there, but obviously this community puts more emphasis on industrial and social aspects of AI, as well as on emerging topics such as AI ethics and AI education and awareness. Notably, this perspective considers AI to be at the same level with robotics.

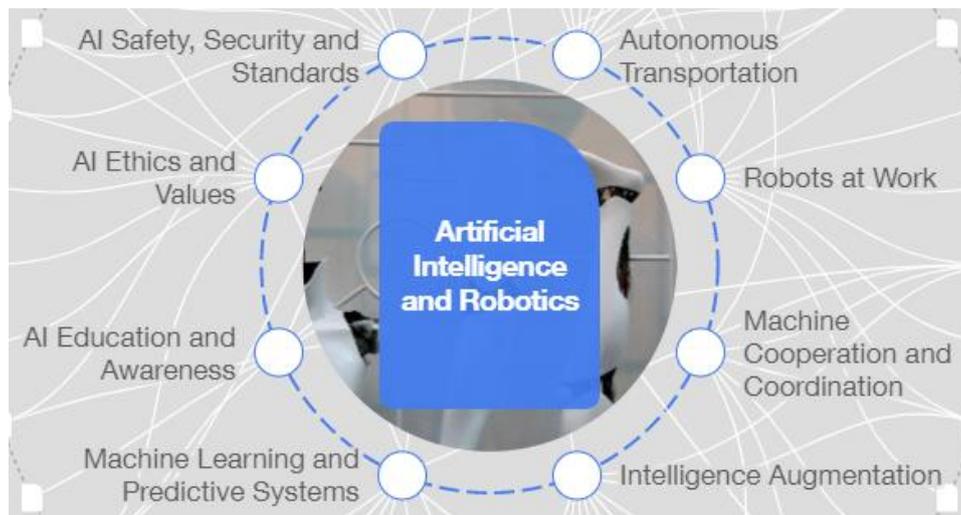


Fig. 4 Current focus in AI as seen by the World Economic Forum (source: <https://intelligence.weforum.org/topics/a1Gb000000pTDREA2?tab=publications>)

4. AI HYPE

The current wave of interest in AI is certainly unsurpassed in the entire history of the field. There have been periods in the past when breakthroughs in AI have received a lot of interest, attention and investments, but then they have been typically followed by periods of disillusionment, AI effect and lack of funding (usually referred to as “AI winters”). This current wave is not only the strongest, but also the longest one. Popular media cover it on a regular basis. Industry, businesses and services invest in AI more than ever before. Year after year universities announce and start new courses and even entire study programs related to AI. Governments open new funding programs and institutions to support further development of AI. Well-known businessmen, investors, entrepreneurs

and even some of the leading AI experts make statements that contribute to the hype (Mark Cuban: “Invest in AI technology or risk becoming ‘a dinosaur’ very soon.”¹⁶; Sundar Pichai: “AI is probably the most important thing humanity has ever worked on”; Koray Kavukcuoglu: “We believe AI will be one of the most powerful enabling technologies ever created – a single invention that could unlock solutions to thousands of problems.”¹⁷; Azamat Abdoullaev: “Whoever creates Real Artificial Intelligence will rule the world.”¹⁸; Andrew Ng: “AI is the new electricity.”¹⁹). Claims like “AI will completely revolutionize our society” are all over the media, and everyone wants to be involved in the technology race [53].

There are several reasons for all the buzz and excitement. The already mentioned technological advancements and largely increased computational power are an important enabler of AI developments, and the available enormous amounts of data come hand in hand with it. Likewise, there really *have* been impressive recent developments that in part justify the hype. For example, some machines can outperform humans in extracting information from images and identifying objects on images [7], [53]. Similarly, in NLP, the latest generative model from OpenAI²⁰, called GPT-3, can generate amazing human-like text on demand [54].

Also, the strategic onlook called Industry 5.0 [55] puts the interaction and collaboration between man and machine right up front and sees AI as one of the major pillars of future industry developments. Promoters envision this important AI trend to make highly automated manufacturing and self-managed supply chains a reality very soon.

Today's technology development leaders like Facebook, Google, Tencent, Amazon, Alibaba etc. all have a great business interest in developing AI-powered systems and applications, and they advertise their efforts. Again, their own success with their AI products is undeniable, and there is no compelling reason why one should believe that they will *not* manage to make next major shifts in that direction.

However, all this interest and attention raises also an important question: Can AI really live up to the hype?

There are opposing opinions, stating that AI has been overhyped and that current AI systems are not very intelligent and thus are very limited. Some already see a decline in the hype, starting from the Gartner hype cycle for AI 2019 that indicates that ML, NLP, DNN and other AI technologies are already on the downward slope of the curve, in the section called the trough of disillusionment [56]. They remind the AI community and the wider public of earlier AI hypes that have crushed by failures (e.g., “the 5th generation of AI”)²¹. They also argue that significant AI results achieved in the past have become part of other disciplines and are no longer considered AI.

¹⁶ <https://yourstory.com/2020/01/ces-2020-mark-cuban-ai-artificial-intelligence-investments-startups>

¹⁷ <https://www.bbc.com/news/technology-51064369>

¹⁸ <https://www.linkedin.com/pulse/global-artificial-intelligence-gai-narrow-ai-applied-ml-dl-abdoullaev/?published=t>

¹⁹ https://www.wipo.int/wipo_magazine/en/2019/03/article_0001.html

²⁰ <https://openai.com/>

²¹ <https://shorturl.at/qHKS3>

Some of the more extreme views in the stream opposing the AI hype even insist that consulting firms deliberately create the fear of missing the AI wave and scare companies into paying for AI projects²². They warn that typical AI applications rarely bring high payoff to companies. AI can be very hard to afford, given the cost of AI specialists and specialized hardware.

Mocking the AI hype comes along the same lines. A famous meme²³ from 2018 makes a parallel between concepts in computing – "then" there were application, program, operating system, script, shell, batch file, service, etc.; in 2010, they have been all replaced by app, app, app,...; in 2018, their names became AI, AI, AI,...

Note, however, that it is not as clear cut (i.e. just promoters vs. opponents) as it might look. The general attitude to AI has changed notably. Once it was not so popular and profitable to start a business with AI. Nowadays, companies proudly wave their AI flags. It has become almost a matter of self-esteem for a company to say that it is not making just ordinary applications, but ones that can learn, talk, perceive objects and so on – much like people – using AI.

When someone makes a pilot study and comes up with results like "In the future, AI will shorten your commute even further via self-driving cars that result in up to 90% fewer accidents, more efficient ride sharing to reduce the number of cars on the road by up to 75%, and smart traffic lights that reduce wait times by 40% and overall travel time by 26%" [1], opponents call it guessing, incomplete, wishful thinking and the like. However, people ask: "How safe are self-driving vehicles? I've heard of an accident caused by malfunction of such a vehicle." Promoters of self-driving vehicles often answer with a counter-question: "How many accidents like that have you heard of?" True, self-driving cars are not that many yet, so the chance of accidents caused by them is still low. If one thinks in terms of percentages/proportions – what are the proportions of the rides that ended up as accidents when a driver was behind the steering wheel, and those that had no driver? An alternative way of thinking about the same problem is: there are no drunk or mad drivers in self-driving cars. Again, the debate is huge, but laymen are very surprised here: some people believe not only that the safety of self-driving cars is not lagging behind that of human-driven cars, but that self-driving cars are *safer*²⁴. They found the grounds for such an opinion in the fact that such vehicles can use much more information than human drivers – information from vehicle-to-vehicle messaging, from ultrasonic and infrared imaging, from automated external traffic-control systems, and so on. Of course, critics will reply that Level-5 (fully automated) self-driving will never be possible because the AI built into self-driving vehicles belongs to a very narrow domain and lacks a wider, human comprehension of the world; thus, the critics say, using a non-humanlike way of achieving intelligence, fully automated and truly intelligent self-driving cars will always "be right around the corner."²⁵

All in all, controversy is already there, but perhaps paradoxically – it only contributes to the hype.

²² <https://www.forbes.com/sites/petercohan/2019/02/15/3-reasons-ai-is-way-overhyped/#31fd61a15a6a>

²³ <https://shorturl.at/mLRUY>

²⁴ <https://qr.ae/TxsdYi>

²⁵ <https://qr.ae/TlyGdY>

5. LIMITATIONS OF WHAT IS CALLED AI TODAY

A good question to ask about the systems that are called AI today is: What *exactly* can these systems do?

A short answer might be: typically, one thing.

For instance, a self-driving car can maybe outperform human drivers in terms of safe driving, communicating with other cars and relevant services to exchange information about road conditions, and even inform the passengers about the route, the driving time, and the like. But it cannot infer how to answer questions like: Who wrote the famous lyrics *Words are flowing out like endless rain into a paper cup?*; or, What does the term *Lonely Planet* stand for? Likewise, after seeing many thousands of images of leopards, a DNN can learn to recognize them with very high accuracy. But it typically breaks when shown an image of a similar animal, like a cheetah, or a lynx. It needs to undergo a time-consuming training process again, to see many thousands of images of cheetahs in order to learn how to recognize them. And the same goes for lynxes. Paradoxically, the process is the same even if it has to learn to recognize something completely different, say a tree. The idea of training another DNN on multiple datasets (e.g., leopards, cheetahs, lynxes and trees) would not work because of feature interference. Even if it worked for a specific multiset, it would face the same problem when possibly adding yet another dataset to the multiset. Efforts to solve this problem do exist (e.g., the proposed multi-modal DL architecture [57] with separate models tuned for each specific dataset in a multiset), but the need for training the resulting DNN again for each new dataset remains.

Actually, the problem is that DNNs are not capable of learning the underlying principles of recognizing similar objects and differentiating them from the starting category of objects.

Just like the fact that AI and ML are not the same things, and that ML is not simply “AI that improves itself” (an idea often found in the popular press), DL is not ML. DL can be superior in learning how to recognize images or natural language, but they are not a magic wand. When it comes to mundane tasks like regression and classification from structured data, like data sourced from a relational database, DL is of little use. In such cases, statistical techniques like gradient boosting [58], e.g. XGBoost [59], are a better choice. Similarly, as Scott E. Fahlman puts it,²⁶ concept detection in NLP using DL works well if a dictionary of words or word patterns representing the concepts of interest is available. Otherwise, traditional symbolic reasoning might be more suitable. On the other hand, symbolic knowledge representation and reasoning techniques are also far from being good in achieving human-level performance in any non-narrow domain, let alone in commonsense reasoning.

ML technology of today is also very limited in terms of generalizing from examples, as well as in terms of learning concepts efficiently and quickly based on a small set of the concept features and on just a few examples. A general problem of most currently popular ML approaches is that they need a lot of data to make statistical inference about possibly existing patterns in the data with acceptable accuracy. The data is typically noisy, and given enough data and enough computing power ML can be successful. However, humans are capable of learning from just a handful of examples and clear data.²⁷ Moreover, a few

²⁶ <https://qr.ae/pNvVSI>

²⁷ <https://qr.ae/TW4h6w>

examples and clear data make it possible for humans to clearly formulate the knowledge the examples convey, to use this knowledge in further reasoning and to explain their reasoning.

Contrary to that, much of ML today works like a black box (with a notable exception of decision trees, which are easily interpretable and explainable). It is especially true for NNs, most notably DNNs. For example, DNNs for image classification can include millions of parameters in their convolutions, ReLU and max pooling layers, which is inherently incomprehensible for humans; explaining how everything works inside such networks is currently an illusion.

Another serious limitation of today's systems called AI is that they are pretty straightforward, which is not typical for intelligent behavior. For example, humans typically drift away in conversations, they change topics, insert jokes and colloquial phrases here and there, and make conversation spontaneous. AI systems don't. True, they can answer questions like "When do I have my next meeting?" and "How long does it take to get from A to B by car?" quite accurately, but they cannot answer any more imaginative questions, like "If Bach was still alive, would he play blues?". In the words of S. Mahadevan, today's AI is designed to answer "What is" questions, but not "What if" questions; the latter "would simply befuddle any AI system".²⁸

Many systems called AI today are also easy to fool. Studies have shown that DNNs are actually very brittle and vulnerable to attacks – making some tiny changes in input images through deliberate adversarial perturbations (like adding some fuzz, noise) [60], even changing only one pixel [61], can lead to a completely wrong classification of the image in a lot of cases. Now, if one thinks of some real-world applications of DNNs, such as self-driving vehicles, such a one-pixel change can be fatal – what if a raindrop "changes this one pixel" in such a way that the car "believes" that a pedestrian is another car? Or, what that one pixel can do if a medical decision is to be made based on a number of images of a tissue? Similarly, an image of a bicycle or a guitar pasted for adversarial purposes over (a part of) an image of a monkey can fool the DNN to classify the animal as a human [62]. The problem here, again, is the black-box nature of DNNs – it is simply difficult to figure out what exactly DNNs are doing inside their hidden layers when they are predicting the class of an input data item, let alone resemblance to how human brains work. Yes, they are always repeating the same algorithmic steps and are making classifications based on some statistics, but humans often have trouble understanding why *such* statistics are dominant. DNNs do *not* model human brains, simply because it is not known how human brain works. More data fed into a DNN can make it more accurate, but not intrinsically human-smart. Also, feeding more data into a DNN cannot account for all possible situations, not even for all possible typical data items; the datasets used contain data from different sources, hence a great deal of repetitive data.

Given all this discussion, one can ask the question: Where is the *intelligence* there?

²⁸ <https://qr.ae/pN2pSZ>

6. REALITY CHECK AND PRACTICAL CHALLENGES

Applying AI to solve practical problems in the real world usually brings up conditions different from those that govern academic research in the field. The understanding of AI (or the lack of such understanding?) in companies and institutions comes from business objectives, which typically command development of technology with more “intelligence”, i.e. with *practical AI* (roughly corresponding to weak AI) and is intentionally limited²⁹. Few companies are interested in developing *general AI* (strong AI), i.e. sentient behavior. Both practical and general AI development require expertise from multiple fields, since “AI is not a single thing”.

6.1. Human-driven AI vs. autonomous AI

Much of practical AI is human-driven. For example, one can see ML as predictive analytics – it creates predictions that inform human decision makers. But all steps in the process – from collecting data into dataset(s) and wrangling with the data to make it suitable for feature engineering, building the model(s), testing them, fitting them and creating predictions – are essentially driven by data engineers / ML engineers. The tools they use do not learn themselves, i.e. to not have a built-in self-improvement logic. Even if such a logic was built in the ML tools, it would still be pre-programmed by human AI specialists. Jeff Bezos calls this human-powered pseudo-AI “AAI” – artificial artificial intelligence.³⁰

In contrast, autonomous AI (general, strong AI) reflects “the very nature of intelligence ... [i.e.] it is self-guided, self-expanding and self-inspired.”³¹ For instance, an ML tool capable of improving its own code, deciding by itself which ML model to use to make predictions, and making different inferences about datasets by itself, would be an autonomous ML tool. To the best of the author’s knowledge, such tools do not exist in practical AI today.

6.2. AI as a marketing term

Sadly, due to the AI hype the label “AI” has largely become a marketing term, and the press and online posts support that situation. It has become “a matter of honor” for companies and institutions to put the label “AI” in their products and profile descriptions, whereas in reality much of the products and activities labeled “AI” are at best applied statistics, business analytics and informed human decision-making.

In marketing, rebranding is a powerful tool. If one looks carefully at the history of terms used to describe parts of research and development often attributed to AI, then they will see that once upon a time there were “pattern matching” and “pattern discovery”. Later on, there came “data mining” and “knowledge discovery” – slightly different, but cultivated on the same soil as their predecessors. Nowadays, all of them are simply rebranded “AI” (or “ML”, or “DL”). From the marketing perspective, it was actually a clever decision: “AI” is catchier, cooler, more appealing and more promising.

Still, just like in any marketing campaign, the reality is different. Today’s dominating weak AI *does* the job in specific narrow application areas, but when compared to general human intelligence – it lives in a galaxy far, far away. As a famous tweet says: When you’re

²⁹ https://www.quora.com/?activity_story=88335643

³⁰ <https://www.wcspeakers.com/speaker/jeff-bezos/>

³¹ <https://medium.com/@ruchika.nanayakkara/ai-is-the-next-virus-42f887a6bec4>

fundraising, it's AI. When you're hiring, it's ML. When you're implementing, it's linear regression.³²

There are also warnings that the hype and hysteria around AI can possibly do *harm* to further AI development [63]. Part of them are based on the fact that the labels "AI" and "ML" are (over)used only to boost sales.³³ As in the tweet mentioned above, "ML" advertises and masks much less popular terms like "regression" and "classification" that would actually describe the essence of ML (and the absence of human-like learning in it) in a more realistic way. However, this "sales pitch" bubble can burst soon, because of the dangers associated with raising expectations too high, without thinking about the real chance of delivering their vision.

Both heavy promoters of AI (often being CEOs in big-name companies, where weak AI is an essential part of their business model) and doom forecasters (predicting massive unemployment due to AI development, existential threat, singularity and even destruction of our civilization – like Stephen Hawking, Elon Musk and Bill Gates, to name but a few) have originally further advertised AI with their statements [63]. However, there is little *evidence* in support of both big promises and big doomsaying. As market research shows, productivity in many countries is slowing down (and not rising) due to automation supported by practical AI, and unemployment is recently at its historical low [63]. Moreover, a 2019 survey conducted by a UK-based investment firm has shown that about 40% of Europe's "AI companies" don't use AI in any way essential to their business [64].

Unfortunately, such facts possibly indicate that the warnings expressed in [63] might be right: once again, as 2019 Gartner curve shows [56], the disillusionment caused by over-advertised but unfulfilled AI promises has started.

6.3. AI seen from different practical perspectives

Different disciplines intersect in what the label "AI" means in the AI community; in a way, as discussed in section 2, it's a catch-all term encompassing subsets of computer science, engineering, statistics, computational linguistics, mathematics, cognitive psychology, neuroscience, philosophy, etc. Even subareas of AI represent intersections of different other disciplines. For example, ML is considered by some as "a rebranding of tools from linear algebra, approximation theory, numerical optimization and statistics."³⁴

Interesting questions here are: What does current AI look like from the perspective of other relevant disciplines? What are the roles of these disciplines in AI? What about industry, employers' expectations and job market? What is the role of AI in a context wider than that of technology development?

6.3.1. The role of statistics

Most ML today heavily depends on statistics; so much, that one can often hear that AI is just statistical fitting (or curve fitting).³⁵ Such statements draw from the fact that, in most ML, conclusions and predictions are made from a large set of training data. In spite of the fact that humans learn differently, from very few examples and making interconnections between different subject areas, experiences and new facts, statistical approaches and NNs in

³² <https://twitter.com/ossia/status/1097804721295773696?lang=en>

³³ <https://qr.ae/pN2r8b>

³⁴ <https://qr.ae/pNsnq3>

³⁵ <https://www.quora.com/When-will-AI-go-beyond-curve-fitting>

ML are dominant in today's AI. S. Mahadevan has put it nicely: "Trying to do ML without knowing statistics is like to trying to build engineering structures without physics."³⁶

In contrast, symbolic AI – by far less popular today than in the past – is often called GOFAI: Good Old-Fashioned AI. It is important to understand that GOFAI, in particular its knowledge representation and reasoning approaches, are not dismissed. Not at all. They bring declarative way of specifying how things should be conducted, strong formalisms of logical reasoning, and also the power of generating explanations. These features can be nicely combined with statistical approaches; for instance, using symbolic approaches rigor can be brought to defining ML pipelines and what exactly they should learn using statistics. In other words, while statistical approaches can process very large, complex data sets, cognitive approaches coming from symbolic AI, like reasoning and problem-solving can bring more human-like flair to AI in order to use AI to its currently possible full potential. ML/Statistical algorithms alone cannot do it; ironically, even some statisticians call ML algorithms "very, very stupid".³⁷ On the other hand, statistical approaches in areas like image recognition and NLP are essential today. It is important to always remember that both statistical and symbolic approaches have their pros and cons.

Note, however, that although much of ML is built on statistics, there is an important difference in approaches between the two: classical statistics always starts from a hypothesis to test, even before the data is collected; ML first collects huge datasets and then applies exploratory statistical analysis in hope to discover some patterns in data and then use them as the model for making predictions.³⁸

It is up to AI course designers at universities to make the role of statistics in AI clear. Unfortunately, it is not always so. In an EDEN Webinar from November 2019 on AI in Higher Education [65], complaints have been put up about courses that have the label "AI" in the title, but are essentially just statistics.

6.3.2. Industry perspective

Google search for "best careers for 2020 and beyond", "best IT career paths for the next decade", "most in-demand IT jobs" and the like, shows controversial results³⁹. A number of Websites ranking such careers does not mention AI and its subareas at all. The "closest" jobs they mention are those of mathematicians, statisticians, operations research analysts, business analyst, market research analysts, marketing specialists (if one assumes that these skills are applied in developing ML models to make analyses). Some Websites rank data analysts, data scientists and data engineers high. Only two such Websites explicitly rank AI architect and robotics engineer high.

A similar search on Indeed.com⁴⁰, driven by queries like "AI", "ML", "AI engineer", "ML engineer", "robotics engineer" and the like, has vaguely reflected the bar graph shown in Fig. 2. However, the "software engineer" query had the number of hits higher by an order

³⁶ <https://qr.ae/pNnJUC>

³⁷ <https://qr.ae/TquNti>

³⁸ <https://qr.ae/pNKFXD>

³⁹ As of Aug. 2020. Only the first few dozens of hits have been surveyed.

⁴⁰ <https://www.indeed.com/>, a popular job announcement and search Website.

of magnitude than the one for “AI engineer”⁴¹. Indeed’s list of 25 best jobs for 2020⁴² includes neither AI nor ML explicitly (“data scientist” is at no. 8, “data engineer” at no. 12).

Related job descriptions reveal the usual AI \equiv ML misconception mentioned in section 3, as well as a frequent vagueness in postings (“Using various techniques, models and algorithms to solve AI problems”, “Applying multiple skills, functional and technical, on AI problems”, “Building prototypes of AI applications”, ...). However, “Strong statistical and math background”, “Programming experience (Java, C/C++, Python, Ruby...)”, “Mathematical and statistical programming experience (R, SAS, SPSS, Python...)” and the like are very frequent accompanying elements in these job announcements as well. In other words, there is much greater demand for job applicants with programming skills and knowledge of statistics than for “pure” AI specialists.

A forum discussion about which undergraduate computer science courses should an aspiring ML engineer take⁴³ lists in the answers AI, ML, probability, statistics, linear algebra, data science, algorithms, and theory of computation, augmented with an introductory course in psychology. Although psychology might look to some as an “outlier” in this list, it actually helps aspiring ML engineers develop a set of skills different from the “core” ones – AI, ML, math, statistics – but also very important in practical work. When ML engineers do not have a good knowledge of the data they have to work with, they have to familiarize with it. In practice, it means attending meetings with the clients and putting a lot of effort in clarifying every single attribute in a dataset.

All these observations should be put in the perspective of *expectations* from both the industry and the job applicants. Actually, many companies expect job applicants to do a lot of data analytics and statistics, rather than DL modeling that is used more frequently in academia⁴⁴. Likewise, most modeling in industry in terms of ML modeling will be traditional modeling, starting from relational databases, not DNN and the like. In addition, due to companies’ expectations, many positions that include ML tasks also comprise programming and software engineering. This often contradicts expectations of job applicants – although all ML includes some programming, it is very different from the programming associated with application development.

Also, most companies use cheap and abundant hardware, which means that the “more data” approach also incurs longer times to train models. Not understanding this important fact and expecting any ML model training to run fast without investing in expensive equipment is a serious misconception. A more-and-more applied strategy to alleviate this problem is to subscribe for cloud-based tools such as AutoML⁴⁵, where training ML models relies on powerful external hardware and software. With tools like that, ML engineers can automate much of the model training, experimentation, fitting and evaluation, getting high-accuracy predictions, but cannot eliminate programming associated with the demanding tasks that precede model building in the ML pipeline – data collection, cleaning and wrangling.

⁴¹ This is probably no wonder at all; in the words of M. Taylor, “Machine Learning is a small part of most projects, and a lot of companies are not going to want to employ a specialist, they are going to expect their software developers to do the job.” (<https://qr.ae/pNkm7b>)

⁴² <https://www.indeed.com/lead/best-jobs-2020>; as of Feb. 2020.

⁴³ <https://qr.ae/pN2tMB>

⁴⁴ Note that there are also different opinions, e.g. <https://qr.ae/pNKKRU>

⁴⁵ <https://cloud.google.com/automl>

From the perspective of an individual company, the workplace roles, the jobs assigned to them and the entire set of business processes and culture should be all tuned well, in order to create new values and make profit. This leaves some room for structured planning and decision-making. A simple tool to use in this process can be a 2×2 matrix with 4 quadrants, defined along the horizontal *Time-to-learn* and vertical *Utility* axes [66]. The quadrants defined this way include *Learn* (high *Utility*, low *Time-to-learn* – the skills and roles that add value for the company quickly), *Plan* (high *Utility*, high *Time-to-learn* – the skills to be acquired only if they are really worth the investment), *Browse* (low *Utility*, low *Time-to-learn* – easy to acquire skills, so stay aware in case their utility increases) and *Ignore* (low *Utility*, high *Time-to-learn* – the company does not have the time for these skills). With this tool, an AI company can simply list the skills it needs (e.g., ML modeling, statistics, data engineering, data collection and wrangling, etc.) and map them onto the four quadrants. The company then typically focuses on the *Learning* quadrant and defines the job roles and positions in a rather straightforward way.

6.3.3. ML engineering and data engineering perspectives

There is some difference between ML engineers and data engineers [67]. ML engineers use programming languages to collect data, clean it, wrangle with it, build and tune ML models and consider alternatives. The languages they typically use include SQL, Python and R. One of the most important and creative activities of ML engineers is feature engineering – what often differentiates successful ML projects from those that fail is the lack of deriving new, useful input features from existing ones.

Data engineers take care of various data sources, formats, storage⁴⁶, infrastructure, scaling and security, and, very importantly, integrating them in applications to make predictions – for example, deploying them in the cloud as microservices [68]. Experience and skills in data ETL (Extract, Transform, Load)⁴⁷ are essential for data engineers, and so is SQL.

These two (often intertwined) job roles make much of “what it really looks like” to work in the area of ML in a company⁴⁸, and is largely different from ML research [69].

Note also that many use the term “data scientist” to encompass ML engineer, data engineer and business analyst roles. This often hinders the real nature of the work done by ML engineers, and some even call this term mislabeling.^{49, 50}

As already mentioned, most of the real work of ML engineers is related to programming. ML model building and tuning takes up to 10-15% of their time (whereas data cleansing and wrangling are about 80% of the job). They work mostly on regression and classification problems, *much* less on DL problems, and their good command of descriptive statistics is understood.

To some, it comes as a surprise that there are usually no entry-level positions for ML engineers and data engineers.⁵¹ But it stops being a surprise when one remembers that, for instance, the ML role assumes knowledge of AI and statistics and a long list of programming and other technical skills. It’s a similar case with the data engineer role.

⁴⁶ <https://qr.ae/pNrDdd>

⁴⁷ <https://qr.ae/pNKF7p>

⁴⁸ <https://qr.ae/pN2yDQ>

⁴⁹ <https://qr.ae/pNypAl>

⁵⁰ <https://qr.ae/pNKKRU>

⁵¹ <https://qr.ae/pN2NUv>

6.3.4. Strategic perspective

No understanding of the current state of affairs in AI can be complete without at least briefly taking into account a more global, strategic perspective.

To this end, the current view is that the strategic leaders in AI are just 9 big companies from China and US [70]: Alibaba (China), Amazon (US), Apple (US), Baidu (China), Facebook (US), Google (US), IBM (US), Microsoft (US) and Tencent (China). Amy Webb, the author of the book [70] specifies: “These companies that are building the frameworks, the custom silicon, it’s their algorithms, it’s their patents. They have the lion’s share of patents in this space. They’re able to attract the top talent. They have the best partnerships with the best universities. It’s these nine companies who are building the rules, systems and business models for the future of artificial intelligence. As a result of that, they have a pretty significant influence on the future of work in everyday life.”⁵²

However, there is a big difference in how these companies work: those from USA are private companies, commercially oriented and with responsibility primarily to their shareholders; those from China, on the other hand, are independent but have to follow the leadership of the government. But in both cases, it is a relatively small group of people that make decisions, and the process is not very transparent.

Application-wise, in USA it is Microsoft that is the leader in defense AI, and Amazon also has a number of contracts with the government related to AI development. Google has pulled out of the defense applications and has focused more on transportation, healthcare and consumer services. When it comes to DL applications, it is Nvidia Corporation that manufactures GPU units that power self-driving vehicles, cloud computing and so on, Deep Instinct is the leader in DL-based cybersecurity, and Microsoft’s cloud computing service, Azure, can run complex DL-driven tools for medical imaging, robotics, NLP etc. In China, AI in transportation has reached an extremely impressive level, and intelligent service robots and drones, neural network chips, and intelligent manufacturing are also among the AI development priorities identified by the Chinese Ministry of Industry and Information Technology.

6.4. Fear of AI vs. benefits of AI

The rapid development of AI and the AI hype have created fear in many people, who seem to believe in the dark predictions mentioned in section 6.2. In a nutshell, the fear is that once intentions, thoughts, human-like behavior and other features of intelligence are coded into programs, machines will become very hard to control and will become inherently dangerous. On the way to this singularity, massive unemployment is almost at sight, in spite of the lack of evidence ([63], [64]) that it looks like that.

Another concern is that the massive data being collected about everything, everywhere, every minute can become a downright threat to privacy and can endanger society by putting control over too many things into hands of governments or other small groups of people. For instance, it has been reported that in China the government has installed over 200 million of surveillance cameras connected with a powerful face-recognition DL system [71]. As a result, each person captured on any of these cameras can be identified and an activity profile is then created for that person. Given the population of China, the technology behind

⁵² <https://www.forbes.com/sites/joemckendrick/2019/04/10/nine-companies-are-shaping-the-future-of-artificial-intelligence/#336612632cf1>

this system is certainly mind-blowing, but the concern is that such an activity profile is then fed into an AI-powered social credit system, meaning that for each person the government calculates a credit score/rating. Those with high scores enjoy benefits in e.g., online purchases, restaurants, hotels and while traveling; those with low scores don't. Sure, companies like Facebook and Google are collecting data about their users and are creating their profiles as well, and it is not clear how they are using these profiles.

A lot of discomfort has also been created by a recent research at MIT, where a DL system called Norman⁵³ has been trained using highly negatively biased data [72]. As a result, images classified in a neutral way by a standard DL image recognition system have been classified by Norman in a scary way. This has raised many concerns, like: "Imagine AI that denies someone a loan because of their gender. Imagine AI that classifies someone as a criminal because of racial prejudice. What's the scariest part of artificial intelligence? How similar it is to us."⁵⁴

Others have rushed to respond quickly, e.g. "There is no reason to give AI control over *goals*. There is only gain to be had in giving it control over *means*... No tool is designed to take over the goals of what it should be used for. Tools don't have their own motives."⁵⁵ They all pull up many examples of "good AI", such as those surveyed in section 1, and their major counter-argument is summarized as "Sometimes those goals, as decided by humans, are dangerous to other humans. But that's not out of control. That's just in the control of a dangerous human."⁵⁶

The largely debated issue that many people will be left jobless and without purpose due to AI-powered automation of many jobs has its reasons. Truck drivers, factory workers, retail and food service assistants are not the only ones to be scared to this end, although their jobs are usually the first ones mentioned in the debates. Stock trading, legal analysis, as well as robotic surgery and medical diagnosis, treatment and care, are often quoted as highly skilled professions where AI will replace humans.

More optimistic views see AI and data revolution as incentives to transform business processes and job roles. The *AI assistant* metaphor is their stronghold – they see AI-driven machines not as competitors for human jobs, but as companions that will do work that they can do better, and will simultaneously let humans focus on things unique to them, such as building relationships, making decisions in complex situations, showing empathy and the like. As G. Warner has nicely put: "Which would you rather have: 1) a human doctor; 2) an AI doctor; or 3) a human doctor using AI?"⁵⁶

Some jobs will certainly cease to exist due to further development of AI – as it has been the case due to different kinds of automation throughout the history of mankind – but some new will be created. In general, many jobs that entail creativity, social interactions, general knowledge, emotional and social intelligence, as well as manual dexterity will thrive; for example, change management specialists, human-computer interaction developers, ML infrastructure maintainers, data curation workers, mental health professionals, etc. An almost "classical" related question is "Will AI replace programmers?" M. Fouts' answer, not without an irony, is: "Every 10 years from 1960 to 1990 at least one major prediction by a prominent AI researcher was "AI will make programmers obsolete in (8-)10 years". 1960

⁵³ <http://norman-ai.mit.edu/>

⁵⁴ <https://qr.ae/pN2KGB>

⁵⁵ <https://qr.ae/TxsB4x>

⁵⁶ <https://qr.ae/pN2Knb>

was 60 years ago and no programmer has ever been replaced by the use of AI software. Nobody has made that prediction since 2000, as far as I know. If AI is ever able to replace programmers, it won't be this century."⁵⁷

In debates on AI pro et contra, there is also a group of people who tend to be neither pessimists nor optimists, but cautious and more realistic, i.e. to see the things from multiple perspectives. Here's a comment coming from that party, in this case with regard to the recently developed GPT-3 natural language generator: "A tool like this has many new uses, both good (from powering better chatbots to helping people code) and bad (from powering better misinformation bots to helping kids cheat on their homework)." [54].

Developing AI that brings benefits to the society is also a concern of governments and political institutions. For instance, European Commission has published a strategic document on development of AI for the benefit of the citizens of EU [73]. The document addresses many opportunities and challenges of AI, but also "a number of potential risks, such as opaque decision-making, gender-based or other kinds of discrimination, intrusion in our private lives or being used for criminal purposes." The guidelines on development of ethical and trustworthy AI [74] have been a precursor to [73]; these guidelines have established a framework for achieving trustworthy AI. The framework has set ethical principles and values for developing AI in Europe, with the idea to foster development of ethical and robust AI. Here "robust" refers to the fact that AI systems can cause unintentional harm, so both technical and social robustness should be addressed when developing an AI system.

6.5. Artificial General Intelligence

Artificial General Intelligence (AGI), also sometimes called General Artificial Intelligence (GAI), has recently proliferated as more-or-less a synonym for strong AI and is used interchangeably with it, as well as with true AI, general AI and real AI (RAI). Conceptually, it is a close approximation of the concept of AI as it was originally envisioned in Mid 1950s – the technology that would be able to do anything that human intelligence can, without human intervention.⁵⁸

Intensive recent discussions about AGI and if it is achievable are largely a side effect of the AI hype. Critics of current AI notice that it is designed only to perform specific tasks, like image recognition and chess playing, tasks that are essentially based on mathematical logic. Fed by huge amounts of data and by pre-programed algorithms, and in some cases equipped by powerful sensory systems (e.g., modern robots and self-driving vehicles), in most mundane applications they do perform well. But if AGI tasks are set as objectives, current approaches simply hit the wall.

An AGI system should also be free of any bias in its behavior, reasoning and actions. This is inherently impossible, if only for the reason of their human designers being biased in many ways (attitudes, objectives, culture and the like)⁵⁹. For instance, Chinese and US AI developers would typically have different views of the AI objectives and purpose). Likewise, AGI is envisioned as observer-independent – also impossible with current technology – whereas

⁵⁷ <https://qr.ae/TI03Vw>

⁵⁸ From Mid 1950s, AI was originally developing that way for approximately 2 decades, before the statistical approach has been initiated in the field.

⁵⁹ <https://qr.ae/pNsn6g>

current AI is observer-dependent.⁶⁰ For example, since human intelligent behavior is typically inseparable from emotions, it is highly unlikely that supporters of animal shelters will react to stray dogs the same way as people who have got bit by such dogs. Last but not least, an essential feature of AGI would be the ability to generalize and then make small variations of the generalized concept or behavior; current AI cannot do it, in spite of some attempts to provide formalisms to do it (e.g., based on description logics [75]). “Throwing larger data sets at faster computers only works for a handful of problems and doesn’t work very well at that... But none of these performances have resulted in a general method that works. Instead, so called data scientists carefully tune data sets used for training, AI companies are caught having humans do what they claim their AI software is doing, and progress has ground nearly to a halt.”⁶¹

Naturally, speculations on the feasibility of AGI have also revived the likewise speculative idea of RAI [37] and have even led to its elaboration into the concepts such as Super Intelligence, Artificial Super Intelligence (ASI), Universal Data Intelligence Framework and the like.⁶² But perhaps more importantly, they have also raised speculations about another AI winter.

There have been two major AI winters in the past (in Early 1970s and Late 1980s / Early 1990s). They have resulted from AI hypes that have preceded them, over-inflated buzz created by popular media and unrealistic promises made by companies and developers. These, in turn, have created extremely high expectations from industry and potential end-users, which have eventually failed to become a reality and have led to the bubble burst effect.

Some base their speculations about another AI winter at sight on making analogies with the previous two. Others⁶³ also look at the Gartner hype cycle for AI 2019 [56], as mentioned in sections 4 and 6.2. Both of these parties express disappointment in current AI not producing commercial results. The hangover is even more obvious from the sheer reality that impressive results in DL and NLP typically come from costly hardware required to train the models with massive data⁶¹ [63]. This especially hits startups, which are beginning to realize that the magic label “AI” alone is not enough to create a ROI. Even big players like Google, Microsoft and OpenAI are beginning to show signs of slowing down the innovation,⁶⁴ since most of their huge ML models still keep mapping input to output, without any reasoning or building world models that AGI supporters demand.

In summary, AGI still remains a myth.

6.6. Challenges

Still, although the hype seems to be declining, there *are* other opportunities and reasonable funding, and there are also intriguing challenges. Some of them are indicated in the Innovation Trigger / On the Rise section at the same Gartner hype cycle for AI 2019 that shows the slight decline of interest in NLP, DL and computer vision [56]. Interestingly, AGI is there, but it is predicted to take more than 10 years before it becomes a reality. Other notable AI technologies on the rise include, e.g.:

⁶⁰ <https://qr.ae/pN2KXE>

⁶¹ <https://qr.ae/TcvCP4>

⁶² <https://shorturl.at/pEFL9>

⁶³ <https://qr.ae/TSTw09>

⁶⁴ <https://qr.ae/TSWZT4>

- Decision intelligence. It is about how to apply ML in organizational decision-making in order to initiate actions with beneficial outcomes. It also applies visualization to help decision-makers quickly grasp cause and effect chains [76].
- Neuromorphic hardware. In this special-purpose hardware, behavior of neurons in human brain is emulated directly in hardware, enabling exceptional and energy-efficient performance during the training of DNNs.⁶⁵
- AI developer kits. This term denotes a set of technologies for straightforward building of AI applications for mobile devices, as well as in the form of Web services.⁶⁶
- AI PaaS (AI platform as a service). Platforms accessible as services for ML developers through a Web-based interface enable developers to build models, use models developed by others, and enjoy the model up- and down-scaling as needed.⁶⁷
- Edge AI. Much of data preprocessing and initial ML can be done by devices used to collect data (e.g., smart speakers), prior to sending data to more powerful computers and servers for further analysis.⁶⁸
- Explainable AI (XAI). In contrast to today's black-box nature of ML, where often even the system designers cannot explain why the model has predicted a specific output, XAI develops with the idea to make the output of an AI system understood by humans [77].
- ...

In addition to these practical development challenges, there is also a number of theoretical challenges that AI still has to take on its path of further expansion. For example, classical questions still without a good theoretical answer are: What exactly is happening inside a NN that makes it possible to train it to recognize images, voices, and so on? Why DL algorithms work? Similarly, how one can infer a suitable number of layers and nodes in a NN? It is still largely a matter of trial and error; there is no theory about it. Likewise, what is the real nature of human vision and can one build a computer vision system based on it, unlike building DL-based image recognition systems where a change of only one pixel can lead to misclassification of the entire image? Along the same lines, can ML work correctly without cleaning noisy data first? Human brain can. In NLP, how to enable semantic understanding of text?

Further on,⁶⁹ instead of just more-or-less accurately mapping a DNN input to output using some (often complicated) transfer function, is it possible to make the network infer some causal knowledge that connects the two? Can a DNN be trained to learn multiple tasks simultaneously? Can it be trained to self-improve over time, possibly in multiple phases, like in the developmental psychology of humans? Ultimately, can it be trained to become self-aware?

These last questions can be tackled in multiple ways. At MIT, researchers have tried to make an AI system evolve on its own, in terms of automatically discovering complete

⁶⁵ <https://www.iis.fraunhofer.de/en/ff/kom/ai/neuromorphic.html>

⁶⁶ <https://www.colocationamerica.com/blog/ai-development-tools>

⁶⁷ <https://geekflare.com/machine-learning-paas/>

⁶⁸ <https://www.digikay.com/en/maker/projects/what-is-edge-ai-machine-learning-iot/4f655838138941138aaad62c170827af>

⁶⁹ <https://qr.ae/pNKDRs>

ML algorithms just using basic mathematical operations as building blocks [78]. Although preliminary results look modest – their evolutionary approach has enabled the system to discover two-layer neural networks trained by backpropagation – it is still extremely promising because of at least two reasons. The first one is the vastness of the search space. While their work has just scratched the surface, it is quite possible that the approach can help discover yet unknown NN algorithms and topologies. The second reason is of at least equal importance: this approach significantly reduces human bias due to a generic search space.

Another group of researchers has made initial progress in developing NNs good for modeling and learning continuous processes (unlike all other NNs, including DNNs, that can model only discrete things, i.e. nothing that transforms continuously over time) [79]. These new NNs are called ODE networks, for Ordinary Differential Equations that parameterize the continuous dynamics of hidden units specified by a neural network. With other NNs, the way training is typically conducted is specifying the number of layers in advance, running the training and then finding how accurate the network is. In contrast, with an ODE network one specifies the target accuracy first, based on which the network configures and trains itself in the most efficient way until it achieves the pre-specified accuracy. The ODE approach is also featured by high memory efficiency. The drawback is that, unlike with other NNs, one cannot tell in the beginning of training how long it will take for an ODE network.

7. CONCLUSIONS?

This is another intentional question mark in a subheading. It is difficult to derive any definite conclusions about AI as a field today, since the only common denominator of so many different views and phenomena is – controversy.

There is still no single, widely adopted and solid definition of what AI is. This is not a surprise, given the fact that there are still a lot of disagreements on what human intelligence is. In spite of that, there seems to be a good deal of agreement about the differences between weak AI and strong AI (AGI), Fig. 5. Still, due to the AI effect, many research results that initially take on the lure of AI, lose that lure over time and become “just technology”. Part of the explanation for that is the fact that virtually all AI today is essentially weak AI, without generalized human cognitive abilities, hence incapable of solving intelligent tasks without human intervention.

It is quite possible that AI effect will not stop until AGI is achieved (if it ever happens). It might also happen that when AGI is achieved the term “AI” will gradually become obsolete and just part of the history of computing.

But until that happens, the reality looks very different. AI cannot do so many things that in the world of humans are taken for granted – e.g., there is still no robot that can implement the moves of an old lady drinking her coffee without spilling the coffee⁷⁰, and no DNN that can recognize the reasons behind a sudden change in a person’s mood. True, advances in technology have accelerated the capture of data and information, and the technology we call ML can usually efficiently analyze this data, build models, and make predictions. But it cannot explain the models and predictions it has made, not at all.

⁷⁰ A brilliant example by A. Kostic, given during an AI-related class at the U. of Belgrade in 2017.

The volume and intensity of the AI hype have created a situation of overselling AI both in industry and in academia. Many businesses declare that they are deploying and/or developing some AI; however, a recent survey has not confirmed it for about 40% of the sample. The offer of AI, ML, DL and similar courses is abundant at universities and at boot camps, and is largely profitable because of people's fear of missing out (despite the employers' reserved opinion about the certificates from such courses). The prophecies of AGI-coming-soon, which the general press is frequently throwing, only contribute to that fear. But few, very few realize some crucial misconceptions about AI, like the one that current AI systems still remain useful in narrow domains. The extreme view is that AI actually doesn't exist.⁷¹

Present

Weak AI
 AI hype
 Domination of DL/DNN, NLP and robotics
 Applications and success in narrow domains
 Data-intensive applications
 Statistical techniques
 Black-box approach, lack of explanation
 Answers to *What is*, not to *What if*
 Brittleness
 Human-driven AI, "AAI"
 AI as a marketing term
 Industry/Research dichotomy
 Overselling AI
 AI race with political implications
 Fear of AI

Future (?)

Strong AI
 AI widely adopted in the society at large
 Human-like AI technology
 AGI at a global level
 Data- and knowledge-intensive applications
 Variety of techniques and approaches
 XAI
 Answers to *What is*, *What if* and more
 Solid, robust systems and applications
 Autonomous AI
 AI as a well understood term
 Better integration of industry and research
 Embedded AI
 Extensive democratization of AI
 Benefits of AI



Fig. 5 A vision of AI

AI has largely become a *metaphor* for data-intensive technology. Is it maybe a sign of a paradigm shift in the field? Long ago, achieving human-level intelligence, or AGI, has been the objective of AI research; supporters of the AGI idea believe that it should remain so. However, AI today seems to be obsessed with data, despite the fact that much of it achieves success only with static data or snapshots of data; but the problem is that data changes over time. Time-series analysis is an approach to tackle this problem, but it is also a data-intensive approach. Things like temporal reasoning, that once have been among the hottest AI topics, seem to be forgotten.

Fortunately, in spite of so many controversies *research* in the broad field of AI is not dead. Researchers (and companies, like Amazon, Baidu, Facebook, Alibaba, OpenAI and Google) always detect and pursue interesting problems at different scales. They often fail to deliver results, but are not afraid to fail – curiosity always prevails over fear (although neither is possible to represent with current AI technology!). Failures indicate the paths not to follow, thus they can still be of some value in the next step.

⁷¹ <https://qr.ae/Tiy96A>

Although nobody knows when and if AGI will be achieved or not, brilliant entrepreneurs and researchers alike keep suggesting how to pursue it. Alan Kay's affirmative attitude about true AI is: "The history of learning how life works is 'very suggestive' that intelligence [can be based on] special organizations of parts that do not at all have to be intelligent into systems that manifest intelligence... From the practical standpoint, it is hard to imagine that solutions will not be more intelligent and reflective than human beings right from the get-go (we are actually terrible thinkers, given what thinking is all about)."⁷² Sridhar Mahadevan seems to share that opinion: "Intelligence emerges from the synergistic interaction of simple entities embedded in complex environments... In this view, we think of intelligence not as an ability innate to a creature, but as a composite of the interactions of the creature with its environment."⁷³

REFERENCES

- [1] D. Faggella, "Everyday Examples of Artificial Intelligence and Machine Learning – Comprehensive Overview," Woburn, MA, Emerj Artificial Intelligence Research, White Paper, 2020.
- [2] T. Stenovec, "Google has gotten incredibly good at predicting traffic – here's how," New York, NY, Business Insider, White Paper, 2015.
- [3] D. Richman, "Uber's machine learning chief says pattern-finding computing fuels ride-hailing giant," Seattle, WA, GeekWire LLC, 2016.
- [4] J. Markoff, "Planes Without Pilots," New York, NY, New York Times, 2015.
- [5] BI Intelligence, "10 million self-driving cars will be on the road by 2020," New York, NY, Business Insider, White Paper, 2015.
- [6] A. Prakash, "Swarm Robotics: New Horizons in Military Research," *Robotics Business Review*, May 2018.
- [7] F. Grimal and J. Jae Sundaram, "Combat Drones: Hives, Swarms, and Autonomous Action?," *J. of Conflict & Security Law*, vol. 23, no. 1, pp. 105–135, Spring 2018.
- [8] L. Huang et al. (Oct. 2011). Adversarial Machine Learning. Presented at AISec'11: 4th ACM Workshop Security and Artificial Intelligence, Chicago, IL. [Online].
- [9] W. Knight, "Military artificial intelligence can be easily and dangerously fooled," *MIT Technology Review*, Oct. 2019.
- [10] N. Mejia, "AI-Based Fraud Detection in Banking – Current Applications and Trends," Woburn, MA, Emerj Artificial Intelligence Research, White Paper, 2020.
- [11] P. Marsden, "Artificial Intelligence Defined: Useful list of popular definitions from business and science," White Paper, 2017.
- [12] S.J. Russell and P. Norvig, *Artificial Intelligence - A Modern Approach*, Third Edition. Boston, MA: Pearson, 2016, Chapter 1, pp. 1–5.
- [13] R.J. Sternberg, "INTELLIGENCE (entry)," in *The Oxford Companion to the Mind*, 1st ed., R.L. Gregory and O.L. Zangwill, Eds., New York, NY, USA: Oxford Univ. Press, 1987, pp. 375–379.
- [14] S. Legg and M. Hutter, "A Collection of Definitions of Intelligence," In Proceedings of the 2007 Conference on Advances in AGI: Concepts, Architectures and Algorithms: Proc. of the AGI Workshop 2006, Jun. 2007, pp. 17–24.
- [15] L.S. Gottfredson, "Mainstream Science on Intelligence: An Editorial with 52 Signatories, History, and Bibliography," *Intelligence*, vol. 24, pp. 13–23, Dec. 1997.
- [16] U. Neisser et al., "Intelligence: Knowns and unknowns," *Amer. Psychologist*, vol. 51, no. 2, 1996, pp. 77–10.
- [17] A. Turing, "Computing machinery and intelligence," *Mind*, vol. 59, no. 236, pp. 433–460, Oct. 1950.
- [18] "Turing test success marks milestone in computing history," U. of Reading press release, Jun. 08, 2014.
- [19] W. Knightley, "Google Duplex: Does it Pass the Turing Test?," Digital Initiative, Harvard Business School, Boston, MA, Nov. 2018.
- [20] "Robots or People: Who's Gonna Rule Tomorrow?," Evergreen, Kyiv, Ukraine.
- [21] J.R. Searle, "Minds, brains, and programs," *Behavioral and Brain Sci.*, vol. 3, no. 3, pp. 417–457, 1980.
- [22] S.E. Fahlman, "How advanced is the most sophisticated example of AI?,"

⁷² <https://qr.ae/pN27DK>

⁷³ <https://qr.ae/pN2pSZ>

- [23] M. Haenlein and A. Kaplan, "A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence," *California Management Review*, vol. 61, no. 4, pp. 5–14, Aug. 2019.
- [24] K. Bailey, "Reframing the 'AI Effect'," San Francisco, CA, Medium Corp., 2016.
- [25] E. Luders et al., "Neuroanatomical correlates of intelligence," *Intelligence*, vol. 37, no. 2, 2009, pp. 156–163.
- [26] A. Nowogrodzki, "The world's strongest MRI machines are pushing human imaging to new limits," *Nature*, vol. 563, no. 7729, pp. 24–26, Nov. 2018.
- [27] S.R. Cox et al., "Structural brain imaging correlates of general intelligence in UK Biobank," *Intelligence*, vol. 76, pp. Sep-Oct. 2019.
- [28] Z. Zheng et al., "A Complete Electron Microscopy Volume of the Brain of Adult *Drosophila melanogaster*," *Cell*, vol. 174, no. 3, pp. 730-743, Jul 19, 2018.
- [29] L.R. Grimm, "Psychology of knowledge representation," *WIREs Cogn. Sci.*, vol. 5, no. 3, pp. 261–270, May-Jun. 2014.
- [30] S. Mahadevan, "How is knowledge representation carried out in the brain?,"
- [31] L. Chang and D.Y. Tsao, "The Code for Facial Identity in the Primate Brain," *Cell*, vol. 169, no. 6, pp. 1013-1028, Jun 2017.
- [32] Leverhulme Centre for the Future of Intelligence, "The Consciousness and Intelligence Project".
- [33] M. Aydede and G. Guzeldere, "Consciousness, intentionality and intelligence: some foundational issues for artificial intelligence," *J. of Experm. & Theor. AI*, vol. 12, no. 3, pp. 263–277, Nov. 2010.
- [34] V. Vinge, "The Coming Technological Singularity: How to Survive in the Post-Human Era," in *Vision-21: Interdisciplinary Science and Engineering in the Era of Cyberspace*, G.A. Landis, ed., NASA Publication CP-10129, pp. 11–22, 1993.
- [35] I.J. Good, "Speculations Concerning the First Ultraintelligent Machine," *Adv. in Computers*, vol. 6, pp. 31–88, 1965.
- [36] R. Kurzweil, *The Singularity is Near*. New York, NY: Viking Books, 2005.
- [37] K. Persianov, "Which company do you think will be the first to create the singularity for artificial intelligence?" .
- [38] M. Brenner, "Why Intelligence might be simpler than we think – Lessons from the Neocortex," San Francisco, CA, Medium Corp., 2019.
- [39] R. Kurzweil, *How to Create Mind?* New York, NY: Viking Books, 2005.
- [40] P. Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. New York, NY: Basic Books, 2015.
- [41] D. Hofstadter, *Gödel, Escher, Bach: An Eternal Golden Braid*. New York, NY: Basic Books, 1979.
- [42] S. Mahadevan, "Imagination Machines: A New Challenge for Artificial Intelligence," Palo Alto, CA, AAAI, 2018.
- [43] T. Mitchell, *Machine Learning*. New York, NY: McGraw Hill, 1997.
- [44] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.
- [45] H. Wang and B. Raj, "On the Origin of Deep Learning," arXiv:1702.07800, 2017.
- [46] A. Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*, 2nd ed. Boston, MA: O'Reilly Media, 2019.
- [47] I. Goodfellow et al., "Generative Adversarial Networks," In Proceedings of the Int. Conf. Neural Inf. Proc. Sys. (NIPS 2014) 2014, pp. 2672–2680.
- [48] T. Young et al., "Recent Trends in Deep Learning Based Natural Language Processing," *IEEE Comp. Intelligence Mag.*, vol. 13, no. 3, pp. 55-75, Aug. 2018.
- [49] H.A. Pierson and M.S. Gashler, "Deep Learning in Robotics: A Review of Recent Research".
- [50] MC.AI, "Fundamentals of Machine Learning (ML), Deep Learning (DL) and Artificial Neural Networks (ANN)," MC.AI, Dec. 11, 2019.
- [51] C. Ramirez, ed., *Advances in Knowledge Representation*. London, UK: IntechOpen Limited, 2012.
- [52] M.K. Bergman, *A Knowledge Representation Practionary: Guidelines Based on Charles Sanders Peirce*. New York, NY: Springer, 2018.
- [53] V. Flovik, "Machine Learning: From hype to real-world applications – How to utilize emerging technologies to drive business value," San Francisco, CA, Medium Corp., TowardsDataScience, Sep 16, 2019.
- [54] W.D. Heaven, "OpenAI's new language generator GPT-3 is shockingly good—and completely mindless," *MIT Technology Review*, Jul. 2020.
- [55] M. Vollmer, "What is Industry 5.0?," Sunnyvale, CA, LinkedIn, August 23, 2018.
- [56] L. Columbus, "What's New in Gartner's Hype Cycle For AI," New York, NY, Forbes Newsletter Group, Sep 25, 2019.
- [57] L. Kaiser et al., "One Model To Learn Them All," arXiv:1706.05137.
- [58] J.H. Friedman, "Greedy function approximation: A gradient boosting machine," *Ann. Statist.* Vol. 29, no. 5, pp. 1189–1232, 2001.

- [59] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," arXiv:1603.02754.
- [60] I.J. Goodfellow, J. Shlens and C. Szegedy, "Explaining and harnessing adversarial examples", arXiv:1412.6572.
- [61] J. Su, D.V. Vargas and S. Kouichi, "One-pixel attack for fooling deep neural networks," arXiv:1710.08864.
- [62] A.L. Yuille and C. Liu, "Limitations of Deep Learning for Vision, and How We Might Fix Them", *The Gradient*, 2019.
- [63] W. Naudé, "AI's current hype and hysteria could set the technology back by decades," *The Conversation*, Jul. 24, 2019.
- [64] W. Knight, "About 40% of Europe's "AI companies" don't use any AI at all," *MIT Technology Review*, Mar. 2019.
- [65] EDEN Network. Artificial Intelligence (AI) in Higher Education. (Nov. 14, 2019).
- [66] C. Littlewood, "Prioritize Which Data Skills Your Company Needs with This 2x2 Matrix," *Harvard Business Rev.*, Oct. 23, 2018.
- [67] M. West, *Acing the Machine Learning Interview*, in press.
- [68] C. Kaiser, "Stop making data scientists manage Kubernetes clusters," San Francisco, CA, Medium Corp., 2019.
- [69] D. Sculley et al., "Hidden Technical Debt in Machine Learning Systems," Corpus ID: 17699480. Accessed Aug. 15, 2020.
- [70] A. Webb, *The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity*. New York City, NY: PublicAffairs, 2019.
- [71] P. Mozur, "Inside China's Dystopian Dreams: A.I., Shame and Lots of Cameras," *New York Times*, Jul. 8, 2018.
- [72] G. Kumar et al., "Scary dark side of artificial intelligence: a perilous contrivance to mankind," *Humanities & Soc. Sci. Rev.*, vol. 7, no. 5, pp. 1097-1103, 2019.
- [73] European Commission, "On Artificial Intelligence - A European approach to excellence and trust," Brussels, COM (2020) 65 final, Feb. 19, 2020. White paper.
- [74] High-Level Expert Group on Artificial Intelligence, "Ethics Guidelines for Trustworthy AI," European Commission, Brussels, Belgium. Apr. 8, 2019.
- [75] A.R. Divroodi et al., "On the possibility of correct concept learning in description logics". *Vietnam J. Comp. Sci.* vol. 5, no. 1, pp. 3–14, 2018.
- [76] C. Byrne, "Why Google defined a new discipline to help humans make decisions," *FastCompany*, Jul. 18, 2018.
- [77] E. Tjoa and C. Guan, "A Survey on Explainable Artificial Intelligence (XAI): Towards Medical XAI," arXiv:1907.07374, 2019.
- [78] E. Real et al., "AutoML-Zero: Evolving Machine Learning Algorithms from Scratch," arXiv:2003.03384, 2020.
- [79] R.T.Q. Chen et al., "Neural Ordinary Differential Equations," arXiv:1806.07366, 2018. Accessed: Aug. 18, 2020.