

Vladan Devedžić*

University of Belgrade, Faculty of Organizational Sciences, Serbia

Quo Vadis, AI?

DOI: 10.7595/management.fon.2021.0001

Abstract:

Research Question: This article reflects on currently popular topics and future development and application of AI. **Motivation:** AI has become a key technology in many areas – from business and services, to science engineering and engineering, natural language processing and news, publishing and writing, human resource management and education, medicine and healthcare, transportation and aviation, government, military and agriculture. **Idea:** As the field advances and expands at an unprecedented pace, it becomes more and more difficult to distinguish between AI and other fields. **Data:** The article is based on a qualitative analysis of relevant literature and direct insights into the field of AI during the last 35 years. **Tools:** The article uses direct observation and content analysis to derive conclusions. **Findings:** Just like any other technology, AI has both its bright and its dark sides. There is a number of expectations, hopes and promises of AI that still remain to be fulfilled, yet new avenues for further development open day after day. **Contribution:** The article attempts to put AI developments in different contexts and shed different lights on AI applications in general, AI applications in business, in particular, and current limitations, trends and challenges in the field.

Keywords: artificial intelligence, machine learning, neural networks, business, applications.

JEL Classification: D80, L63, L86

1. Introduction

It is very difficult, if not impossible, to define Artificial Intelligence (AI). The field is huge, it also overlaps with other fields and uses approaches and techniques from other fields. AI focuses on mimicking parts of intelligent human behaviour in machines, and human intelligence itself is extremely difficult to define – it is simply much too complex to allow for capturing in a single definition. As R.J. Sternberg (1987) has put it, “viewed narrowly, there seems to be almost as many definitions of intelligence as there are experts asked to define it”. All the aspects of human intelligence – thinking, reasoning, abstracting, generalizing, planning, and many more – are also extremely complex; there is currently no scientific understanding of the nature of these processes.

For the purpose of further discussion, a characterization of intelligence offered by Henrik H. Lund from the University of Southern Denmark is adopted here: intelligence is human “capability of creating a variety of behaviours, while complying with the givens of the system/environment.”²

In an attempt to understand the current state-of-the-art in AI at the time of writing this article, the article briefly presents several examples of systems labelled AI and puts the current AI hype in the context of Gartner Hype Cycle curves for AI over the recent years. There is also a discussion on what businesses and managers should bear in mind when talking about AI and making decisions accordingly. The author's reflections on recent developments in AI follow, mentioning several important aspects that usually do not get much attention, converging into a set of indications of possible answers to the question put in the title of the article.

¹ Corresponding author: Vladan Devedžić, e-mail: devedzic@fon.rs

² Copied from one of H.H. Lund's slides during his presentation that the author has attended long ago.

2. Examples of AI

One of the best-known examples of AI is that of *personal assistants*, such as Amazon Alexa (Amazon Alexa, 2020) and Google Assistant (Google Assistant, 2020). These devices provide intelligent personal assistance, in terms of helping human users find the information they are looking for, add/retrieve events to/from their calendars, or giving directions when the users ask for them. The communication between the users and such devices is conducted in (a limited form of) the natural language.

Self-driving vehicles, such as those manufactured by the Tesla company (Tesla, 2020), collect relevant information from the environment and are capable of taking passengers and goods around without being driven by human drivers.

When consumers want to watch a movie on Netflix (<https://www.netflix.com/>), Netflix uses its *predictive, AI-based technology* (Kaplan & Haenlein 2019; Aggarwal, Goswami, Hooda, Chakravarty & Kar 2020) to analyze the user's previous choices and recommend other films to watch. Similarly, Amazon (<https://www.amazon.com/>) predicts what the users are interested to purchase (Zhang & Pennacchiotti, 2013; Zeng, Cao, Chen & Li 2019) based on their online behaviours.

Email services such as Gmail (<https://mail.google.com/>) successfully filter spam (Dada, Bassi, Chiroma, Adetunmbi & Ajibuwa, 2019) and suggest the users how to create replies to messages, pointing to which emails they have not replied to or have ignored. Navigating with Google Maps (<https://www.google.com/maps/>) is easy, since Google Maps intelligently calculates the route (Hu, Wu, Cho, & Tseng 2020) taking into account the traffic conditions and road construction in order to recommend a faster route.

Automated image analysis (Kazeminia et al., 2020) and *object detection/recognition* (Liu et al., 2019) have become popular as well, and there are more and more services that the users can read their images into and the services will recognize the scene, parts of the scene and details on the images. Likewise, *speech-to-text services* (Myers, 2020) enable voice-based input of search queries when using search engines, as spoken input gets automatically converted to text (the search string).

Applications and services that are called AI today are typically based on training the system with large volumes of data in order to build a model that then accurately classifies input and predicts output when presented with new data from the same category. The accuracy is never 100%, however, it has already got impressive.

At the time of writing this article (Sept. 2020), the technology that enables many of these applications is *deep learning* (Goodfellow, Bengio, & Courville, 2016). It uses neural networks with complex topography in order to successfully recognize different features of sound, images and other forms of input data and integrate them into accurate higher-level descriptions that allow the network to produce correct output (such as classification, prediction, and the like). Deep learning is a subset of *machine learning*, a very important part of what is today called AI.

3. Gartner Hype Cycle for AI

The well-known Gartner Hype Cycle curve, Figure 1, provides a view of how a certain technology or application will evolve over time. Gartner Inc. publishes such curves each year, for different technologies, based on their complex methodology of analyzing technological trends.

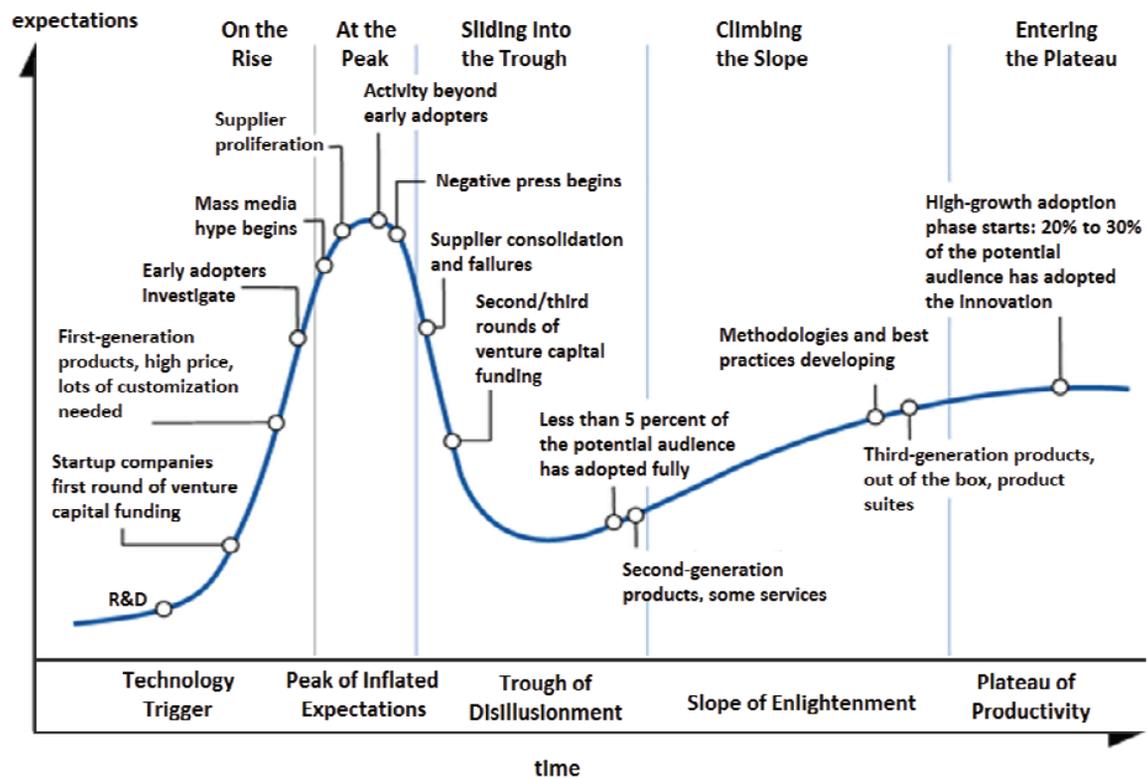


Figure 1: Gartner Hype Cycle curve, general (source: <https://commons.wikimedia.org/wiki/File:Hype-Cycle-General.png>, marked for reuse under Creative Commons license)

The leftmost section of the curve (Technology Trigger) typically starts with some interesting research findings that investors, start-ups and venture capitalists see as promising in terms of generating some profit if the findings are built into products. This is followed by media promotion and advertising, and, if it generates enough attention among the general public, it can create expectations that are often not quite realistic (the Peak of Inflated Expectations section). Failing to fulfill these expectations, the technology/product(s) gradually slide into the Trough of Disillusionment, where the general press is not interested in promoting it further. Still, some adopters and some investors remain interested, which is the enabler for the second-generation products that make the curve rise again, though at a slower pace and with more cautious expectations (Slope of Enlightenment). As more experience with the new products is gained and best practices are gradually developed, the curve reaches The Plateau of Productivity, with larger adoption (but still lower than it has been expected at the peak of the hype).

If one follows Gartner Hype Cycle curves for AI in recent years (hint: Google search for “Gartner Hype Cycle AI”, and see Google Images for the results), some trends emerge there:

- deep learning is at the peak of inflated expectations for at least three years; it has started to slightly decline in 2020, but it is still there;
- traditional machine learning is declining, as it was predicted (Liu 2018) and is now sliding into the Trough of Disillusionment;
- even Natural Language Processing (NLP), in spite of the recent breakthroughs, seems to be slightly declining;
- Artificial General Intelligence (AGI) (Rohrer, 2011) is on the rise; it has been there for quite some time, but the 2020 trend shows some ascension; AGI refers to the endeavours of developing intelligent machines capable of performing intelligent tasks at the human level, rather than focusing on narrow aspects of intelligence without understanding the real nature of intelligence;
- AI cloud services, or AI PaaS (platform as a service) has been on the rise for a number of recent years, and has entered the peak of popularity in 2020; AI PaaS is about putting a number of AI services (such as deep learning and natural language translation) in the cloud, making them run on powerful hardware that speeds up the training processes, and enabling end-users to access them through the Internet;
- Edge AI – delegating AI processing to peripheral (edge) devices, rather than centralizing them on cloud services – is also getting popular;

- Small Data (Olson, Wyner & Berk 2018) movement is on the rise – instead of using massive data volumes to train machine learning (typically deep learning) models, how about learning from just a handful of data, as humans typically do?

4. AI, Business and Managers

Managers and businesses should be aware of some important facts when it comes to AI today:

- the AI hype is evident everywhere; it's undeniable;
- still, it is not clear whether AI can live up to the hype – there have always been promises coming from AI researchers, early adopters of AI technology and general press, many of which have failed to deliver;
- what is often labeled 'AI' today is essentially not much more than brilliantly applied statistics; no true intelligence or true learning in the human-like sense is evident in today's AI applications; in other words, no AGI exists today;
- one of the reasons why many deny any intelligence in AI systems is that they are still human-driven, and not autonomous; likewise, there is a lot of (human-generated) bias in AI systems as they are designed by humans who are inherently biased;
- yet, when surveyed for their opinion on whether AI can replace managers in companies or not (Schawbel, 2019), many respondents in different countries provided an affirmative answer; they only believed that AI was still inferior to humans when it comes to soft-skill management – showing empathy, coaching, promoting work culture and the like – for most other typical managerial roles, like financial and human resource management, the respondents believed that AI-powered robots can do better than humans;
- AI has become an important *marketing* term – in many businesses and services it has become for companies "a matter of survival" on the market to advertise themselves as AI companies, or at least as companies that use AI in their business; however, a 2019 survey conducted by a UK-based investment firm has shown that about 40% of Europe's "AI companies" do not use AI in any way essential to their business (Knight, 2019);
- when hiring AI specialists to work for a company today, it is essential to realize that the set of skills needed includes a lot of database analysis and administration, a good command of SQL, experience in programming, and a thorough understanding of data security; it is estimated that these activities consume 90% of AI specialists at workplaces, whereas for only 10% of their time they work on building and evaluating AI models;
- model democratization (Rao, 2020) is an important factor to take into account – with the rise of AI PaaS, such as AutoML from Google, everyone can build and evaluate their models online, running cloud-based services that perform great (compared to human analysts) and companies do not have to care about providing and maintaining expensive hardware for model building and running.

Akter, Michael, Uddin, McCarthy & Rahman (2020) give an overview of digital business transformation with AI. Still, much of their overview can be subsumed by the third bullet point from above. In their own words: "AI can only work through learning from a vast amount of existing data. For example, Airbus used their AI system to examine a production problem, calculate a vast amount of data, ... *the first step is to use big data, the second is to apply analytics, and the third is prediction*".

Similarly, Marr (2019) lists examples of how industry giants from USA and China use AI in their businesses. His text reads, e.g.: "Artificial intelligence (AI) is integral in Alibaba's daily operations and is used to *predict* what customers might want to buy... [Amazon] They collect a lot of data about each person's buying habits and have such confidence in how the data they collect help them recommend items to their customers and now predict what they need even before they need it by using predictive analytics... Tencent acquires huge amounts of information and insights about its customers that it processes and leverages to the company's advantage...". And so on. Now, go back to the last sentence of the previous paragraph and read the highlighted part; do you see a common pattern here?

And then Faggella (2020) quotes an example from General Electric ("... take vast amounts of information recorded over time to develop its forecasts") and other companies. DigitalSilk (2020) covers even more examples, the one from Coca Cola being typical ("... big data gives the intelligence to cap social media, mobile applications cloud computing and eCommerce off – giving companies like Coca Cola a toolset the way they approach IT and business... This was backed by the more recent statement by Coca Cola global

director of digital innovation, Greg Chambers, that confessed 'AI is the foundation for everything we do', as they 'create intelligent experiences.' "). See again the third and the sixth bullet point from above and the things become crystal clear.

5. Memories, Dreams, Reflections³

This section is inspired by the author's ~35-year R&D experience in the field of AI, as well as by the concerns about further developments in AI that model democratization might bring, as expressed in (Rao, 2020).

The author has started his study of AI at pre-Internet times, when AI was understood much as the science and engineering of developing intelligent systems that, true, *did* focus on narrow aspects of intelligent behaviour as they do today, but also *were* always rooted in the understanding of intelligent behaviour as a very complex one, in the sense expressed by Henrik Lund (see the Introduction section). Moreover, the author remembers previous AI hypes, but also *AI winters* – periods in the history of the field characterized by the lack of funding and considerable decrease in interest in AI by both investors/industry and academia.

The author is also perfectly aware of two phenomena that often shadow the results and achievements in developing intelligent systems. One of them is called *AI effect* (Haenlein and Kaplan, 2019). It refers to reactions to any kind of AI technology once it becomes widely understood and built into commercial products – that is the moment when an AI finding/result that has initially generated a lot of attention stops being that alluring and becomes an ordinary and massively used product that people take for granted. Moreover, at that moment, this result typically also loses its 'AI' label and becomes 'just technology'. For example, long ago expert systems were developing under the AI umbrella, and then they each have become a technology per se. They have even almost lost their original names – now people more often refer to them as rule-based systems. The same goes for chatbots – nobody calls them AI chatbots anymore; they are just chatbots.

The other phenomenon is called *the frame problem* (McCarthy and Hayes, 1969; Dennett, 1984). Put simply, the frame problem means that it is very difficult, if not impossible, to specify all possible knowledge that characterizes the environment in which an AI system works. Typically, the environment changes in response to the system's actions, and practical AI systems are designed with limited sets of beliefs of what can be changed in the environment in response to their actions. Also, all actions that the system takes are specified by what exactly they do change when they happen and assuming that everything else in the environment will remain unchanged. However, the problem is how to create a manageable default set of rational and plausible assumptions of what does not change in the environment; it is tightly coupled with commonsense knowledge, which is inherently difficult to model and represent in practical AI systems.

Referring again to Figure 1, there were several periods in the development of AI when a technological trigger was followed by hype and inflated expectations. Technological triggers also trigger imagination, and if media join the game then many people start dreaming about things that are simply not possible. Dreamers often forget that:

- AGI still does not exist;
- the current understanding of the phenomenon of intelligence is still far from being sufficient to develop human-level AI systems;
- human-level AI systems should be based on important principles of human intelligence that enable that variety of behaviors that Lund points to;
- unfortunately, these important principles and the very nature of intelligence is still unknown.

Some developments are not possible due to the limitations of current technology. In other cases, they are not possible because of the frame problem. As a result, even nowadays all AI developments result in very narrowly focused practical systems that definitely lack a variety of behaviours and, due to AI effect, few people would call them intelligent (although the 'AI' label is still attached to these developments).

It is also an open question whether a large variety of behaviours is possible without the autonomy of AI systems, in terms of adapting to the environment themselves without human intervention. Self-driving cars are a good example of *partial* autonomy, but they are not built on true, human-level intelligence.

³ The title of this section is borrowed from the title of the autobiography of C.G. Jung, published in 1961. However, the author was tempted to entitle this section *Is there any AI out there?*

In business, it is important to be aware of such facts. Consulting firms often try to scare businesses and their CEOs with stories about how costly it can be to miss the current AI wave; this is how they get their profit. But investing in what is called AI can be even more costly. What is called AI nowadays requires a lot of knowledge of math and statistics, and powerful and expensive hardware is often required for the purpose of training AI models. With AI PaaS, things seem to be much more affordable, but... Here's a quote from (Rao, 2020):

"When company leaders talk about democratizing artificial intelligence (AI), it is easy to imagine what they have in mind. The more people with access to the raw materials of the knowledge, tools, and data required to build an AI system, the more innovations are bound to emerge. Efficiency improves and engagement increases. Faced with a shortage of technical talent Microsoft, Amazon, and Google have all released premade drag-and-drop or no-code AI tools that allow people to integrate AI into applications without needing to know how to build machine learning (ML) models."

However, this is exactly where problems can start. As Rao (2020) carefully notices: "Bias can lead to dangerous consequences, and they can become evident only after the product has been released... The key for company leaders is to avoid getting carried away by the hype." This can be interpreted as the fact that bias can lead to malicious developments and abuse of democratized models. It can also lead to a misuse of PaaS, as well as to overselling of possibly poor-quality PaaS.

It is also important for company leaders to understand the fact that much of what is nowadays sold as AI is actually applied statistics and a lot of curve fitting. It is a good idea for them to take a glance at two famous memes⁴⁵ that illustrate this point. AI has become so much overhyped and the label 'AI' so much oversold that everything is now called 'AI'.

Then, it is also good to pause and think: Is intelligence just model building and curve fitting? True, the business world today is simply obsessed with data. More data, more data, more data ... Why? To better fit the curves and call that AI? What is wrong with calling it statistics? (Except that maybe AI is a more catchy word?) The famous slogan that data is the fuel of modern economy is probably true today, but there is a question with no consensual answer yet: Is curve fitting by building models with more and more data sufficient to call it machine learning (or, even more generally, AI)?

This brings up several other important questions. What does *learning* have to do with it? Do, for example, new employees in a company have to see millions of company's reports in order to figure out (build model, learn) how to write such reports? After all, do children learn what cats are by being shown millions of cats? Is 'machine learning' a misnomer for an actually great application of statistics to nowadays extremely important part of business – data analytics?

Rao (2020) further warns: "Yet as we curate data, algorithms, and models ... the risk of misinterpreting them and applying them in the wrong context increases significantly. The danger of systemic misuse of models will rise... But who is accessing these (democratized) tools and models? Have those users been appropriately trained – not just in the tool, but also in the underlying concepts? ... Business users are typically casual users; they have received extensive training neither in the statistical and mathematical concepts underlying the models nor in the specific processes required to build models."

To illustrate how important this extensive training can be, note that neural networks (including deep ones) are very easy to fool – Su et al. have shown (2019) that *changing only one pixel* in the input image can make the network classify the image in a completely wrong way. Moreover, training deep learning models is an extremely complex process and it is practically impossible to know what exactly is happening inside the network during training. As a consequence, showing the network an image of an object that it is trained to recognize (say, a monkey) but on an unusual background, can also lead to an erroneous classification (say, the monkey can be classified as a human (Yuille and Liu, 2019)).

⁴ <https://devrant.com/rants/1494050/machine-learning>

⁵ <https://shorturl.at/mLRUY>

This is to say that only thoroughly trained people should take responsibility in developing and deploying systems that we usually label AI – those who do understand that, for example, a single drop of rain on the camera of a combat drone can change not one, but several pixels of its input image, resulting in a wrong selection of the target to release the payload to.

Unfortunately, casual users of the so-called AI technology very often do not understand such potential dangers. It is not unusual to hear suggestions like "Put some AI in there..." in designing a software product or app, as if it was "Move that image on the screen 2cm to the left". Under the surface, this indicates a serious lack of not only training, but also *understanding* of what true AI means (or at least what it should be about) beyond using it as a marketing term.

In reality, this lack of understanding spans much further than just casual users in companies. Even the most prominent researchers in AI, in neurosciences, in psychology, and in all other disciplines, still do not have a good understanding of the very nature of intelligence. This is why all of the models in any subdiscipline of AI and in any AI technology are built only on some creative, but *unverified assumptions* about intelligence. To illustrate this point, note that a great success in studying how the brain of an insect works (an *insect*, not a human!) has been achieved in 2018, when researchers have successfully used electron microscopy to map synaptic connectivity in the brain of fruit flies (*Drosophila melanogaster*) in an unbiased way (Zheng et al., 2018). However, it is still far away from understanding how the brain of that insect works, let alone being able to synthesize that brain – such an achievement is not even at sight. And now compare the brain of a fruit fly – the size of a poppy seed – to the human brain.

Is Jeff Bezos right when he calls today's human-powered pseudo-AI "AAI" – artificial artificial intelligence (Schwartz, 2019)?

Conclusion

So, Quo Vadis, AI?

It is very hard to say. On the one hand, the hype is still on and it is a fact that no previous AI hype has lived this long (it is since more than a decade that the last AI hype has started). One has to admit that the results achieved during this last period of AI development have been more impressive than ever before (see section 2). Many AI systems are now trained on large volumes of data, and today data are available in tons – in companies, in government institutions, all over the Internet. Advances in hardware development have contributed to the acceleration of processing these large volumes of data.

On the other hand, critics say that no AI system today is capable of exhibiting true, human-like intelligence and that what most people label 'AI' is actually just technology that is powerful and impressive, but not smart. Thus, many predict that the AI bubble will burst soon and that the AI community might even experience another AI winter – cease of funding for AI research and development, due to the disillusionment and disappointment in technology that has been promising much, but eventually fails to deliver.

As guidelines in trying to decide where AI is heading today and how to drive business decisions to this end, it might be helpful to:

- watch for trends; Gartner Hype Cycle curves for AI are a good source of information for businesses and managers, but keeping an eye on well-reputed AI community Websites like AI Topics (<https://aitopics.org/search>), Figure 2, is probably even better for spotting emerging trends early:

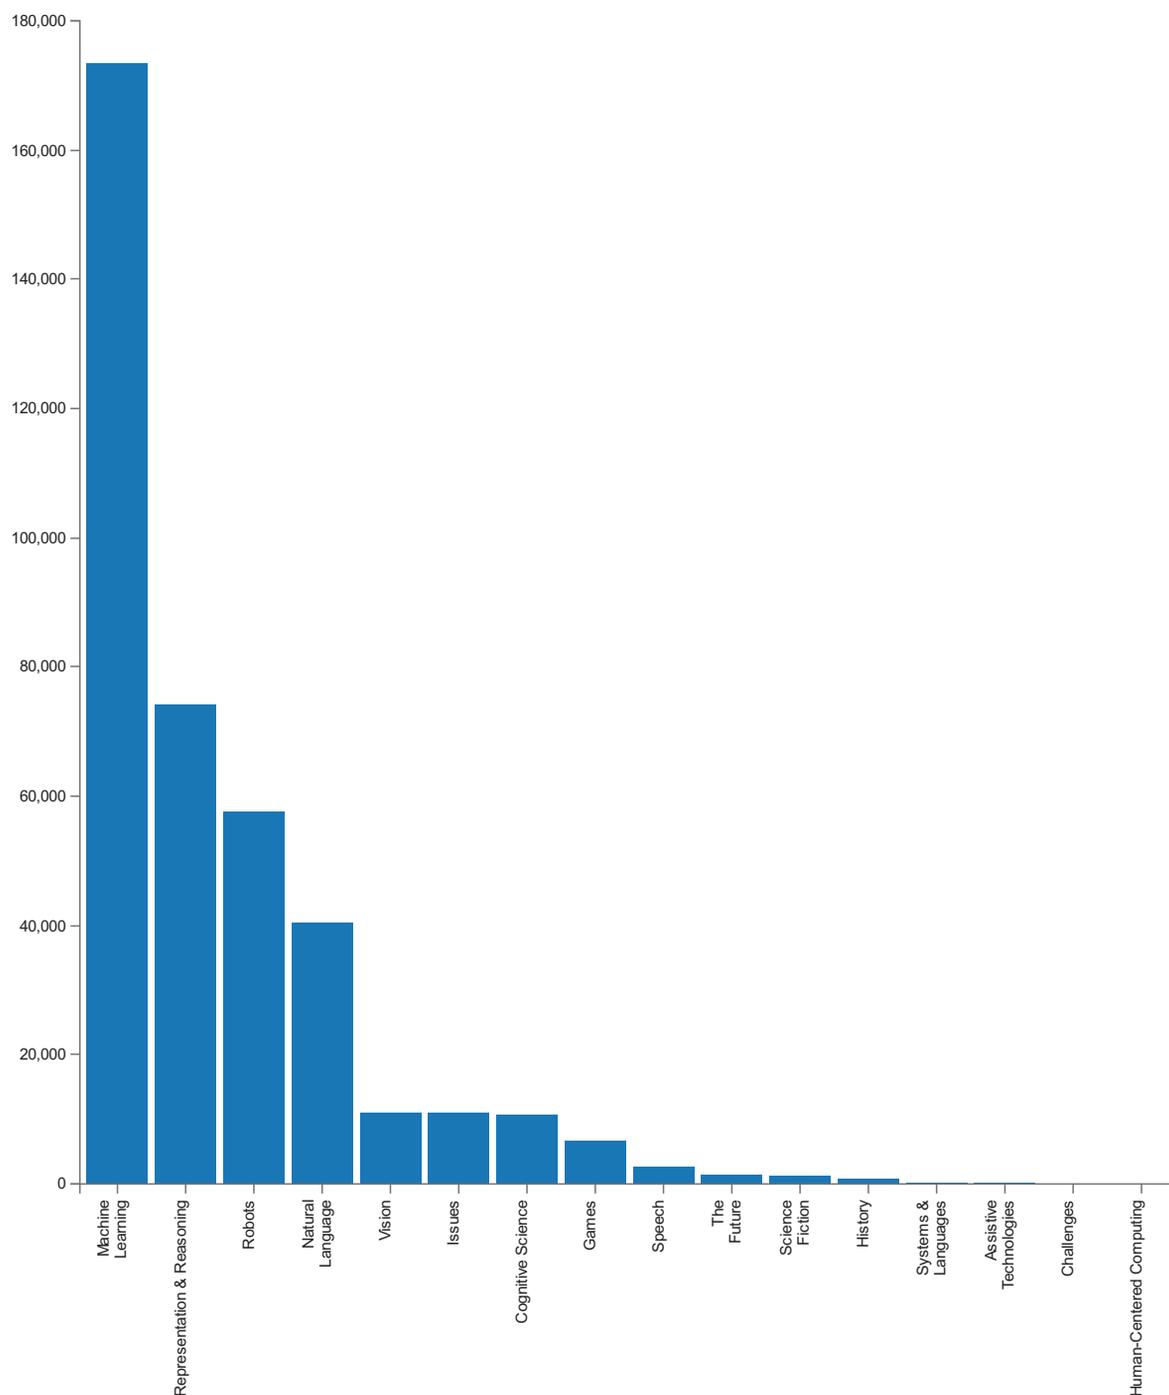


Figure 2: Current AI trends (source: AI Topics, Sep 2020, <https://aitopics.org/explore/classifications?filters=taxnodes:Technology%7CInformation%20Technology%7CArtificial%20Intelligence&from-taxnode=Technology%7CInformation%20Technology%7CArtificial%20Intelligence>)

- remember the context; Figure 3 shows a non-exhaustive list of traditional AI topics, almost always discussed in books on AI – according to many, *all* these topics constitute AI as a field, not only a handful of those exposed on the AI Topics Website (Figure 2);
- bear in mind that, although it is not clear from Figures 2 and 3, human intelligence (capabilities like search, knowledge representation, reasoning, problem solving etc.) is largely based on *heuristics*, which current AI systems largely lack;

Traditional AI topics

Search	Robotics	Agents
Representation	NLP	Neural networks
Reasoning	Vision	Fuzzy systems
Problem solving	Speech recognition	Genetic algorithms
Planning	Machine learning	Swarm intelligence
Scheduling	Pattern recognition	Fractals

Figure 3: Topics studied in AI since its early days

- be aware of the (sad?) fact that the term 'AI' is often abused; in his famous tweet, Quincy Larson mockingly notes (<https://twitter.com/ossia/status/1097804721295773696?lang=en>):

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

Acknowledgments

Parts of this paper have been presented as invited lecture at the XVII International Symposium Symorg 2020 –“Business and Artificial Intelligence”. The paper has been recommended for further extension and review in the Management: Journal of Sustainable Business and Management Solutions in Emerging Economies.

REFERENCES

- [1] Aggarwal, S., Goswami, D., Hooda, M., Chakravarty, A., & Kar, A. (2020). Recommendation Systems for Interactive Multimedia Entertainment. In *Data Visualization and Knowledge Engineering* (pp. 23-48). Springer, Cham.
- [2] Akter, S., Michael, K., Uddin, M. R., McCarthy, G., & Rahman, M. (2020). Transforming business using digital innovations: the application of AI, blockchain, cloud and data analytics. *Annals of Operations Research*, 1-33.
- [3] Amazon Alexa (entry) (2020). Wikipedia. Retrieved from https://en.wikipedia.org/wiki/Amazon_Alexa
- [4] Dada, E. G., Bassi, J. S., Chiroma, H., Adetunmbi, A. O., & Ajibuwa, O. E. (2019). Machine learning for email spam filtering: review, approaches and open research problems. *Heliyon*, 5(6), e01802.
- [5] Dennett, D.C. (1984). Cognitive wheels: the frame problem of AI. In C. Hookway (ed.), *Minds, Machines and Evolution*, Cambridge University Press, 129–150.
- [6] DigitalSilk (2020). Artificial Intelligence In Business: Benefits, Impact & 5 Industry Examples. Retrieved from https://www.digitalsilk.com/artificial-intelligence-in-business#5_Examples_Of_Artificial_Intelligence_In_Business_For_Different_Segments
- [7] Faggella, D. (2020, Mar 14). 6 Examples of AI in Business Intelligence Applications. Retrieved from <https://emerj.com/ai-sector-overviews/ai-in-business-intelligence-applications/>
- [8] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. Cambridge, MA: MIT Press, 2016. Retrieved from <http://www.deeplearningbook.org>
- [9] Google, LLC (2020). Google Assistant. Retrieved from <https://assistant.google.com/>
- [10] Haenlein, M., & Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. *California management review*, 61(4), 5-14. M. Haenlein and A. Kaplan (2019). 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.
- [11] Hu, W. C., Wu, H. T., Cho, H. H., & Tseng, F. H. (2020). Optimal Route Planning System for Logistics Vehicles Based on Artificial Intelligence. *Journal of Internet Technology*, 21(3), 757-764.
- [12] Kaplan, A., & Haenlein, M. (2019). Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15-25.

- [13] Kazemina, S., Baur, C., Kuijper, A., van Ginneken, B., Navab, N., Albarqouni, S., & Mukhopadhyay, A. (2020). GANs for medical image analysis. *Artificial Intelligence in Medicine*, 101938. S. Kazemina et al. (2020). GANs for medical image analysis. *Artificial Intelligence in Medicine*, Vol.109, Sep. 2020. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0933365719311510?via%3Dihub>
- [14] Knight, W. (2019, Mar). About 40% of Europe's "AI companies" don't use any AI at all. *MIT Technology Review*. Retrieved from <https://www.technologyreview.com/2019/03/05/65990/about-40-of-europes-ai-companies-dont-actually-use-any-ai-at-all/>
- [15] Liu, L., Ouyang, W., Wang, X., Fieguth, P., Chen, J., Liu, X., & Pietikäinen, M. (2020). Deep learning for generic object detection: A survey. *International journal of computer vision*, 128(2), 261-318.
- [16] Liu, V. X. (2018). Towards the 'Plateau of Productivity': enhancing the value of machine learning in critical care. *Critical care medicine*, 46(7), 1196.
- [17] Marr, B. (2019, Dec 09). The 10 Best Examples of How Companies Use Artificial Intelligence in Practice. Retrieved from <https://www.forbes.com/sites/bernardmarr/2019/12/09/the-10-best-examples-of-how-companies-use-artificial-intelligence-in-practice/?sh=6dd7f00b7978>
- [18] McCarthy, J., & Hayes, P. J. J. McCarthy and P.J. Hayes (1969). Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence* 4:463–502.
- [19] Myers, E. (2020). The 2020 Speech to Text Report is Here (New Research Report). Retrieved from <https://www.rev.com/blog/speech-to-text-new-research-report>
- [20] Olson, M., Wyner, A., & Berk, R. (2018). Modern neural networks generalize on small data sets. In *Advances in Neural Information Processing Systems* (pp. 3619-3628).
- [21] Rao, A. (2020). Democratizing artificial intelligence is a double-edged sword. *Tech & Innovation*, June 15, 2020. Retrieved from <https://www.strategy-business.com/article/Democratizing-artificial-intelligence-is-a-double-edged-sword?gko=ffdcd>
- [22] Rohrer, B. (2010). Accelerating progress in Artificial General Intelligence: Choosing a benchmark for natural world interaction. *Journal of Artificial General Intelligence*, 2(1), 1-28. B. Rohrer (2011). Accelerating progress in Artificial General Intelligence: Choosing a benchmark for natural world interaction. *Journal of Artificial General Intelligence*, Vol.2, No.1. Retrieved from <https://content.sciendo.com/view/journals/jagi/2/1/article-p1.xml>
- [23] Schawbel, D. (2019, Nov 15). How artificial intelligence is redefining the role of manager. Retrieved from <https://www.weforum.org/agenda/2019/11/how-artificial-intelligence-is-redefining-the-role-of-manager>
- [24] Schwartz, O. (2019, Apr 22). Untold History of AI: How Amazon's Mechanical Turkers Got Squeezed Inside the Machine. *IEEE Spectrum*. Retrieved from <https://spectrum.ieee.org/tech-talk/tech-history/dawn-of-electronics/untold-history-of-ai-mechanical-turk-revisited-tkkt>
- [25] Sternberg, R.J. (1987). INTELLIGENCE (entry). *The Oxford Companion to the Mind*, 1st ed., R.L. Gregory and O.L. Zangwill, Eds., New York, NY, USA: Oxford Univ. Press, 375–379.
- [26] Su, J., Vargas, D. V., & Sakurai, K. (2019). One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, 23(5), 828-841. J. Su, D.V. Vargas and S. Kouichi (2019). One-pixel attack for fooling deep neural networks. arXiv:1710.08864. Accessed: Aug. 12, 2020. Retrieved from <https://arxiv.org/pdf/1710.08864.pdf>
- [27] Tesla Inc. (2020). Tesla Model S. Retrieved from <https://www.tesla.com/>
- [28] Yuille A.L. and C. Liu C. (2019). Limitations of Deep Learning for Vision, and How We Might Fix Them", *The Gradient*, 2019. Accessed: Aug. 12, 2020. Retrieved from <https://thegradients.pub/the-limitations-of-visual-deep-learning-and-how-we-might-fix-them/>
- [29] Zeng, M., Cao, H., Chen, M., & Li, Y. (2019). User behaviour modeling, recommendations, and purchase prediction during shopping festivals. *Electronic Markets*, 29(2), 263-274.
- [30] Zhang, Y., & Pennacchiotti, M. (2013, May). Predicting purchase behaviors from social media. In *Proceedings of the 22nd international conference on World Wide Web* (pp. 1521-1532).
- [31] Zheng, Z., Lauritzen, J. S., Perlman, E., Robinson, C. G., Nichols, M., Milkie, D., ... & Calle-Schuler, S. A. (2018). A complete electron microscopy volume of the brain of adult *Drosophila melanogaster*. *Cell*, 174(3), 730-743. Z. Zheng et al. (2018). A Complete Electron Microscopy Volume of the Brain of Adult *Drosophila melanogaster*. *Cell*, Vol. 174, No. 3, pp. 730-743. Retrieved from <https://www.cell.com/action/showPdf?pii=S0092-8674%2818%2930787-6>

Received: 2020-11-12

Revisions requested: 2020-12-17

Revised: 2020-12-28

Accepted: 2020-12-28



About the Author

Vladan Devedžić

University of Belgrade, Faculty of Organizational Sciences, Serbia
devedzic@fon.rs



Vladan B. Devedzic, PhD, is a Professor of Software Engineering at FOS – Faculty of Organizational Sciences, University of Belgrade, Serbia. He also used to teach several computer science courses at the School of Electrical Engineering, University of Belgrade, as well as at the Military Academy of Serbia (ex Yugoslav Military Academy).

His long-term professional objective is to bring close together ideas from the broad fields of intelligent systems and software engineering. His current professional and research interests include software engineering, programming education, intelligent software systems and technology-enhanced learning (TEL). He is the founder and the chair of the GOOD OLD AI research network. Further details about his work and research can be found at <http://devedzic.fon.bg.ac.rs/>

