



Artificial Intelligence: A Clarification of Misconceptions, Myths and Desired Status

Frank Emmert-Streib^{1*}, Olli Yli-Harja^{2,3,4} and Matthias Dehmer^{5,6,7}

¹Predictive Society and Data Analytics Lab, Faculty of Information Technology and Communication Sciences, Tampere University, Tampere, Finland, ²Institute of Biosciences and Medical Technology, Tampere, Finland, ³Computational System Biology, Faculty of Medicine and Health Technology, Tampere University, Finland, ⁴Institute for Systems Biology, Seattle, WA, United States, ⁵Department of Mechatronics and Biomedical Computer Science, UMIT, Hall in Tyrol, IL, Austria, ⁶Department of Computer Science, Swiss Distance University of Applied Sciences, Brig, Switzerland, ⁷College of Artificial Intelligence, Nankai University, Tianjin, China

The field artificial intelligence (AI) was founded over 65 years ago. Starting with great hopes and ambitious goals the field progressed through various stages of popularity and has recently undergone a revival through the introduction of deep neural networks. Some problems of AI are that, so far, neither the “intelligence” nor the goals of AI are formally defined causing confusion when comparing AI to other fields. In this paper, we present a perspective on the desired and current status of AI in relation to machine learning and statistics and clarify common misconceptions and myths. Our discussion is intended to lift the veil of vagueness surrounding AI to reveal its true countenance.

Keywords: artificial intelligence, artificial general intelligence, machine learning, statistics, data science, deep neural networks, data mining, pattern recognition

OPEN ACCESS

Edited by:

Thomas Hartung,
Johns Hopkins University,
United States

Reviewed by:

Thomas Luechtefeld,
Toxtrack LLC, United States
Lihua Feng,
Central South University, China

*Correspondence:

Frank Emmert-Streib
v@bio-complexity.com

Specialty section:

This article was submitted to
Machine Learning and Artificial
Intelligence,
a section of the journal
Frontiers in Artificial Intelligence

Received: 03 January 2020

Accepted: 12 October 2020

Published: 23 December 2020

Citation:

Emmert-Streib F, Yli-Harja O and
Dehmer M (2020) Artificial Intelligence:
A Clarification of Misconceptions,
Myths and Desired Status.
Front. Artif. Intell. 3:524339.
doi: 10.3389/frai.2020.524339

1. INTRODUCTION

Artificial intelligence (AI) has a long tradition going back many decades. The name *artificial intelligence* was coined by McCarthy at the Dartmouth conference in 1956 that started a concerted endeavor of this research field which continues to date (McCarthy et al., 2006). The initial focus of AI was on *symbolic models and reasoning* followed by the first wave of *neural networks* (NN) and *expert systems* (ES) (Rosenblatt, 1957; Newel and Simon, 1976; Crevier, 1993). The field experienced a severe setback when Minsky and Papert demonstrated problems with perceptrons in learning non-linear separable functions, e.g., the exclusive OR (XOR) (Minsky and Papert, 1969). This significantly affected the progression of AI in the following years especially in neural networks. However, in the 1980s neural networks made a comeback through invention of the *back-propagation algorithm* (Rumelhart et al., 1986). Later in the 1990s research about *intelligent agents* garnered broad interest (Wooldridge and Jennings, 1995) exploring for instance coupled effects of perceptions and actions (Wolpert and Kawato, 1998; Emmert-Streib, 2003). Finally, in the early 2000s big data became available and led to another revival of neural networks in the form of deep neural networks (DNN) (Hochreiter and Schmidhuber, 1997; Hinton et al., 2006; O’Leary, 2013; LeCun et al., 2015).

During these years, AI has achieved great success in many different fields including robotics, speech recognition, facial recognition, healthcare, and finance (Bahrammirzaee, 2010; Brooks, 1991; Krizhevsky et al., 2012; Hochreiter and Schmidhuber, 1997; Thrun, 2002; Yu et al., 2018). Importantly, these problems do not all fall within one field, e.g., computer science, but span a multitude of disciplines including psychology, neuroscience, economy, and medicine. Given the breath of AI applications and the variety of different methods used, it is no surprise that seemingly

simple questions, e.g., regarding the aims and goals of AI are obscured especially for those scientists who did follow the field since its inception. For this reason, in this paper, we discuss the desired and current status of AI regarding its definition and provide a clarification for the discrepancy. Specifically, we provide a perspective on AI in relation to machine learning and statistics and clarify common misconceptions and myths.

Our paper is organized as follows. In the next section, we discuss the desired and current status of artificial intelligence including the definition of “intelligence” and strong AI. Then we clarify frequently encountered misconceptions about AI. Finally, we discuss characteristics of methods from artificial intelligence in relation to machine learning and statistics. The paper is completed with concluding remarks.

2. WHAT IS ARTIFICIAL INTELLIGENCE?

We begin our discussion by clarifying the meaning of artificial intelligence itself. We start by presenting attempts to define “intelligence” followed by informal characterizations of AI as the former problem is currently unresolvable.

2.1. Defining “Intelligence” in Artificial Intelligence

From the name “artificial intelligence” it seems obvious that AI deals with an artificial—not natural—form of intelligence. Hence, defining ‘intelligence’ in a precise way will tell us what AI is about. Unfortunately, currently, there is no such definition available that would be generally accepted by the community. For an extensive discussion of the difficulties encountered when attempting to provide such a definition see, e.g., (Legg and Hutter, 2007; Wang, 2019).

Despite the lack of such a generally accepted definition, there are various attempts. For instance, a recent formal measure has been suggested by (Legg and Hutter, 2007). Interestingly, the authors start from several informal definitions of *human* intelligence to define machine intelligence formally. The resulting measure is given by

$$\text{universal intelligence of agent } \pi = Y(\pi) = \sum_{\mu \in E} 2^{-K(\mu)} V_{\mu}^{\pi}. \quad (1)$$

Here π is an agent, K the Kolmogorov complexity function, E the set of all environments, μ one particular environment, $2^{-K(\mu)}$ the algorithmic probability distribution over an environment and V_{μ}^{π} a value function. Overall, $Y(\pi)$ is called the universal intelligence of agent π (Legg and Hutter, 2007). Informally, **Eq 1** gives a measure for intelligence as the ability of an agent to achieve goals in a wide range of environments (Legg and Hutter, 2007).

A general problem with the definition given in **Eq 1** is that its form is rather cumbersome and unintuitive, and its exact practical evaluation is not possible because the Kolmogorov complexity function K is not computable but requires an approximation. A further problem is to perform intelligence tests because, e.g., a Turing test (Turing, 1950) is insufficient (Legg and Hutter, 2007),

for instance, because an agent could appear intelligent without actually being intelligent (Block, 1981).

A good summary of the problem in defining “intelligence” and AI is given in (Winston and Brown, 1984), who state that “Defining intelligence usually takes a semester-long struggle, and even after that I am not sure we ever get a definition really nailed down. But operationally speaking, we want to make machines smart.” In summary, there is currently neither a generally accepted definition of “intelligence” nor tests that could be used to identify “intelligence” reliably.

In spite of this lack of a general definition of “intelligence,” there is a philosophical separation of AI systems based on this notion. The so called weak AI hypothesis states that “machines could act as if they were intelligent” whereas the strong AI hypothesis asserts “that machines that do so are actually thinking (not just simulating thinking)” (Russell and Norvig, 2016). The latter in particular is very controversial and an argument against a strong AI is the Chinese room (Searle, 2008). We would like to note that strong AI has recently been rebranded as *artificial general intelligence* (AGI) (Goertzel and Pennachin, 2007; Yampolskiy and Fox, 2012).

2.2. Informal Characterizations of Artificial Intelligence

Since there is no generally accepted definition of “intelligence” AI has been characterized informally from its beginnings. For instance, in (Winston and Brown, 1984) it is stated that “The primary goal of Artificial Intelligence is to make machines smarter. The secondary goals of Artificial Intelligence are to understand what intelligence is (the Nobel laureate purpose) and to make machines more useful (the entrepreneurial purpose)”. Kurzweil noted that artificial intelligence is “The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil et al., 1990). Furthermore, Feigenbaum (Feigenbaum, 1963) said “artificial intelligence research is concerned with constructing machines (usually programs for general-purpose computers) which exhibit behavior such that, if it were observed in human activity, we would deign to label the behavior ‘intelligent’.” The latter reminds one of a Turing test of intelligence and that a measure for intelligence is connected to such a test; see our discussion in the last section.

Feigenbaum further specifies that “One group of researchers is concerned with simulating human information-processing activity, with the quest for precise psychological theories of human cognitive activity” and “A second group of researchers is concerned with evoking intelligent behavior from machines whether or not the information processes employed have anything to do with plausible human cognitive mechanisms” (Feigenbaum, 1963). Similar distinctions have been made in (Simon, 1969; Pomerol, 1997). Interestingly, the first point addresses a natural—not artificial—form of cognition showing that some scientists even cross the boundary from artificial to biological phenomena.

From this follows, that from its beginnings, AI had high aspirations focusing on ultimate goals centered around

intelligent and *smart* behavior rather than on simple questions as represented, e.g., by classification or regression problems as discussed in statistics or machine learning. This also means that AI is not explicitly data-focused but assumes the availability of data which would allow the studying of such high-hanging questions. This relates also, e.g., to probabilistic or symbolic approaches (Koenig and Simmons, 1994; Hoehndorf and Queralt-Rosinach, 2017). Importantly, this is in contrast to data science which places data at the center of the investigation and develops estimation techniques for extracting the optimum of information contained in data set(s) possibly by applying more than one method (Emmert-Streib and Dehmer, 2019a).

2.3. Current Status

From the above discussion, it seems fair to assert that we neither have a generally accepted, formal (mathematical) definition of “intelligence” nor do we have one succinct informal definition of AI that would go beyond its obvious meaning. Instead, there are many different characterizations and opinions about what AI should be (Wang, 2006).

Given this deficiency it is not surprising that there are many misconceptions and misunderstandings about AI in general. In the following section, we discuss some of these.

3. COMMON MISCONCEPTIONS AND MYTHS

In this section, we discuss some frequently encountered misconceptions about AI and clarify some false assumptions.

AI aims to explain how the brain works. No, because brains occur only in living (biological) beings and not in artificial machines. Instead, fields studying the molecular biological mechanisms of natural brains are neuroscience and neurobiology. Whether AI research can contribute to this question in some way is unclear but so far no breakthrough contribution has been made. Nevertheless, it is unquestionable that AI research was inspired by neurobiology from its very beginnings (Fan et al., 2020) and one prominent example for this is Hebbian learning (Hebb, 1949) or extensions thereof (Emmert-Streib, 2006).

AI methods work similar to the brain. No, this is not true; even if the most popular methods of AI are called neural networks which are inspired by biological brains. Importantly, despite the name “neural network” such models do not present physiological neural models because neither the model of a neuron nor the connectivity between the neurons in neural networks is biologically plausible nor realistic. That means neither the connectivity structure of convolutional neural networks nor that of deep feedforward neural networks or other deep learning architectures are biologically realistic. In contrast, a physiological model of a biological neuron is the Hodgkin-Huxley model (Hodgkin and Huxley, 1952) or the FitzHugh-Nagumo model (Nagumo et al., 1962) and the large-scale connectivity of the brain is to date largely unknown.

Methods from AI have a different purpose as methods from machine learning or statistics. No, the general purpose of all

methods from these fields is to analyze data. However, each field introduced different methods with different underlying philosophies. Specifically, the philosophy of AI is to aim at ultimate goals, which are possibly unrealistic, rather than to answer simple questions. As a note, we would like to remark that any manipulation of data stored in a computer, is a form of data analysis. Interestingly, this is even true for agent-based systems, e.g., robotics, which incrementally gather data via the interaction with an environment. Kaplan and Haenlein phrased this nicely as “a system’s ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation”, when defining AI (Kaplan and Haenlein, 2019).

AI is a technology. No, AI is a methodology. That means the methods behind AI are (mathematical) learning algorithms that adjust the parameters of methods via learning rules. However, when implementing AI methods certain problems may require an optimization of the method in combination with computer hardware, e.g., by using a GPU, in order to improve the computation time it takes to execute a task. The latter combination may give the impression that AI is a technology but by downscaling a problem one can always reduce the hardware requirements, demonstrating the principle workings of a method, potentially for toy examples. Importantly, in the above argument we emphasized the intellectual component of AI. It is clear that AI cannot be done with a pencil and piece of paper, hence, a computer is always required and a computer is a form of technology. However, the intellectual component of AI is not the computer itself but the software implementing learning rules.

AI makes computers think. From a scientific point of view, no, because similar to the problems of defining “intelligence” there is currently no definition of “thinking”. Also thinking is in general associated with humans who are biological beings rather than artificial machines. In general, this point is related to the goals of strong AI and the counter argument by Searle (Searle, 2008).

Why does AI appear more mythical than machine learning or statistics? Considering the fact that those fields serve a similar purpose (see above) this is indeed strange. However, we think that the reason therefore is twofold. First, the vague definition of AI leaves much room for guesswork and wishful thinking which can be populated by a wide range of philosophical considerations. Second, the high aspirations of AI enable speculations about ultimate or futuristic goals like “making machines think” or “making machines human-like”.

Making machines behave like humans is optimal. At first, this sounds reasonable but let us consider an example. Suppose there is a group of people and the task is to classify handwritten numbers. This is a difficult problem because the hand writing can be difficult to read. For this reason, one cannot expect that all people will achieve the optimal score, but some people perform better than others. Hence, the behavior of every human is not optimal compared to the maximal score or even the best performing human. Also, if we give the same group of people, a number of different tasks to solve, then it is unlikely that the same person will always perform best. Altogether, it does not make sense to make a computer behave like humans because most people do not perform optimally, regardless of what task we

TABLE 1 | List of popular, core artificial intelligence, machine learning and statistics methods representing characteristic models of those fields.

Model	Application	Refs.
Neural networks (NN)	Function approximation, classification	Rosenblatt (1957), Schmidhuber (2015)
Expert system (ES)	Knowledge-based decisions	Hayes-Roth et al. (1983)
Hidden markov models (HMM)	Sequential symbol processing	Rabiner (1989)
Bayesian networks (BN)	Uncertain reasoning	Pearl (1988), Scutari (2010)
Reinforcement learning (RL)	Decision planning	Kaelbling et al. (1996), Sutton and Barto (1998)
Support vector machines (SVM)	Classification	Vapnik (1995), Schölkopf and Smola (2002)
Adaptive boosting (AB)	Classification	Freund and Schapire (1997)
XGBoost (XGB)	Classification	Chen and Guestrin (2016)
Locally linear embedding (LLE)	Nonlinear dimensionality reduction	Roweis and Saul (2000)
Random forests (RF)	Classification	Breiman (2001)
Linear regression (LR)	Regression	Weisberg (2005), Emmert-Streib and Dehmer (2019b)
Logistic regression (LogR)	Classification	Kleinbaum et al. (2002)
Generalized linear models (GLM)	Regression	Nelder and Wedderburn (1972), Dunn and Smyth (2018)
Statistical hypothesis testing (SHT)	Hypothesis testing	Sheskin (2004), Emmert-Streib and Dehmer (2019d)
Cox proportional hazard model (CPHM)	Survival analysis	Cox (1972), Kleinbaum and Klein (2005)

consider. So, what it actually means is to make a computer perform like the best performing human. For one task this may actually mimic the behavior of one human, however, for several tasks this will correspond to the behavior of a different human for every task. Hence, such a *super human* does not exist. That means if a machine can solve more than one task it does not make sense to compare it to one human because such a person does not exist. Hence, the goal is to make machines behave like an ideal super human.

When will the ultimate goals of AI be reached? Over the years there have been a number of predictions. For instance, Simon predicted in 1965 that “Machines will be capable, within twenty years, of doing any work a man can do” (Simon, 1965), Minsky stated in 1967 that “Within a generation . . . the problem of creating artificial intelligence will substantially be solved” (Minsky, 1967) and Kurzweil predicted in 2005 that strong AI, which he calls singularity, will be realized by 2045 (Kurzweil, 2005). Obviously, the former two predictions turned out to be wrong and the latter one remains in the future. However, predictions about undefined entities are vague (see our discussion about intelligence above) and cannot be evaluated systematically. Nevertheless, it is unquestionable that methods from AI make a continuous contribution to many areas in science and industry.

From the above discussion one realizes that metaphors are frequently used when talking about AI but those are not meant to be understood in a precise way but more as a motivation or stimulation. The origin of this might be related to the community behind AI which is considerably different from the more mathematics oriented communities in statistics or machine learning.

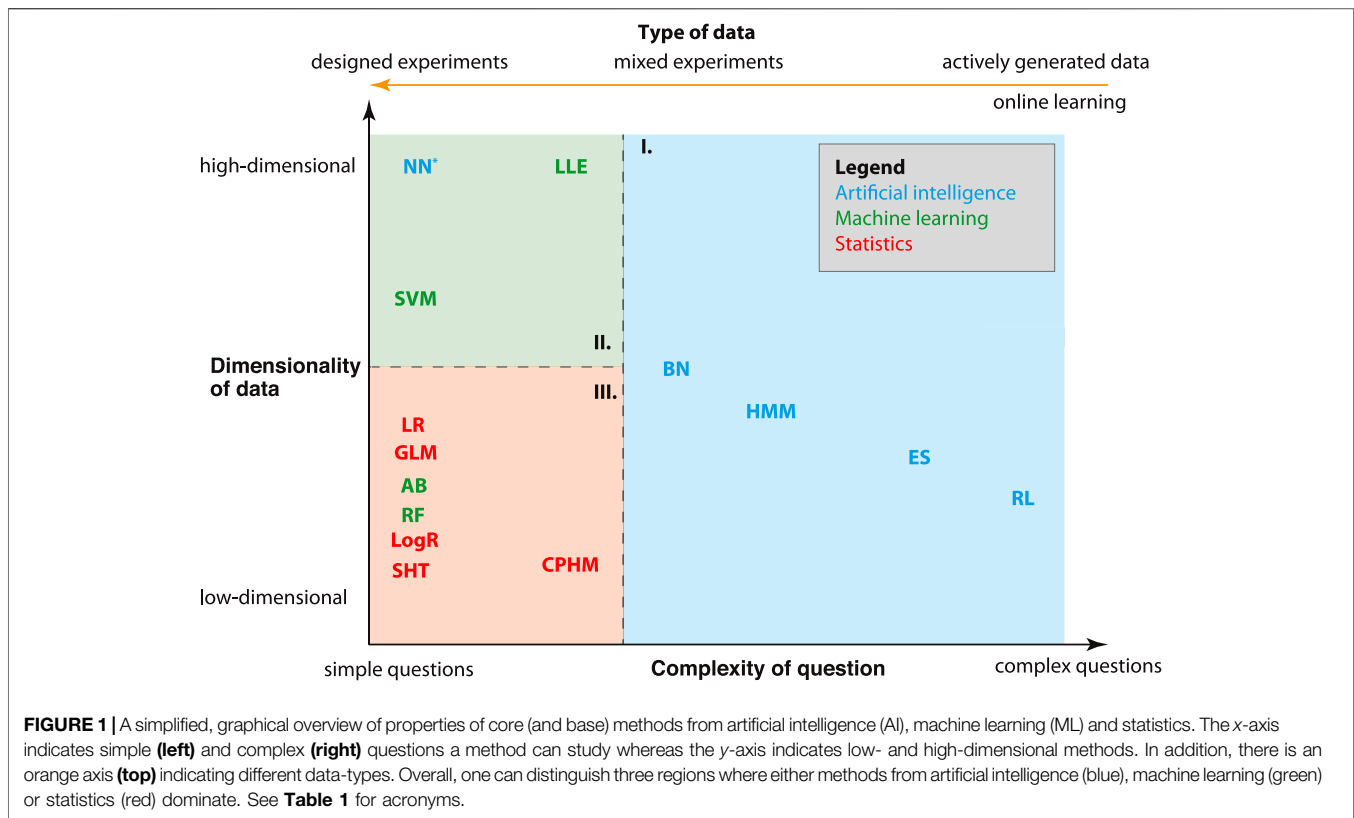
4. DISCUSSION

In the previous sections, we discussed various aspects of AI and their limitations. Now we aim for a general overview of the relations between methods in artificial intelligence, machine learning and statistics. In **Table 1** we show a list of core

methods from artificial intelligence, machine learning and statistics. Here “core” refers to methods that can be considered as characteristic for a field, e.g., hypothesis testing for statistics, support vector machines for machine learning or neural networks for artificial intelligence. Each of these methods has attributes with respect to the capabilities of the methods. In the following, we consider three such attributes as the most important; 1) the complexity of questions to be studied, 2) the dimensionality of data to be processed and 3) the type of data that can be analyzed. In **Figure 1**, we show a simplified, graphical overview of these properties (the acronyms are given in **Table 1**). We would like to highlight that these distinctions present our own, idealized perspective shared by many. However, alternative views and perspectives are possible.

In general, there are many properties of methods one can use for a distinction. However, we start by focusing on only two such features. Specifically, the x -axis in **Figure 1** indicates the question-type that can be addressed by a method from simple (left-hand side) to complex (right-hand side) questions, whereas the y -axis indicates the input dimensionality of the data from low-to-high-dimensional. Here the dimensionality of the data corresponds to the length of a feature vector used as the input for an analysis method which is different to the number of samples which gives the total number of different feature vectors. Overall, in **Figure 1** one can distinguish three regions where methods from artificial intelligence (blue), machine learning (green), or statistics (red) dominate. Interestingly, before the introduction of deep learning neural networks, region II. was entirely dominated by machine learning methods. For this reason we added a star to neural networks (NN) to indicate it as modern AI method. As one can see, methods from statistics are generally characterized by simple questions that can be studied in low-dimensional settings. Here by “simple” we do not mean “boring” or “uninteresting” but rather “specific” or “well defined”. Hence, from **Figure 1** one can conclude that AI tends to address complex questions that do not fit well into a conventional framework, e.g., as represented by statistics. The only exception is neural networks.

For most of the methods shown in **Table 1** exist extensions to the “base” method. For instance, a classical statistical hypothesis



test is conducted just once. However, modern problems in genomics or the social sciences require the testing of thousands or millions of hypotheses. For this reason multiple testing corrections have been introduced (Farcomeni, 2008; Emmert-Streib and Dehmer, 2019c). Similar extensions can be found for most other methods, e.g., regression. However, when considering only the original core methods one obtains a simplified categorization for the domains of AI, ML and statistics, which can be summarized as follows:

- Traditional domain of artificial intelligence \Rightarrow Complex questions
- Traditional domain of machine learning \Rightarrow High-dimensional data
- Traditional domain of statistics \Rightarrow Simple questions

In **Figure 1**, we added one additional axis (orange) on top of the figure indicating different types of data. In contrast to the axes for the question-type and the dimensionality of input data, the scale of this axis is categorical, which means there is no smooth transition between the corresponding categories. Using this axis (feature) as an additional perspective, one can see that machine learning as well as statistics methods require data from designed experiments. This form of experiment corresponds to the conventional type of experiments in physics or biology, where the measurements follow a predefined plan called an experimental design. In contrast, AI methods frequently use actively generated data [also known as online learning (Hoi

et al., 2018)] which become available in a sequential order. An example for this data type is the data a robot generates by exploring its environment or data generated by moves in games (Mnih et al., 2013).

We think it is important to emphasize that (neither) methods from AI (nor from machine learning or statistics) can be mathematically derived from a common, underlying methodological framework but they have been introduced separately and independently. In contrast, physical theories, e.g., about statistical mechanics or quantum mechanics, can be derived from a Hamiltonian formalism or alternatively from Fisher Information (Frieden and Frieden, 1998; Goldstein et al., 2013).

Maybe the most interesting insight from **Figure 1** is that the current most successful AI methods, namely neural networks, do not address complex questions but simple ones (e.g., classification or regression) for high-dimensional data (Emmert-Streib et al., 2020). This is notable because it goes counter the tradition of AI taking on novel and complex problems. Also considering the current interest in futuristic problems, e.g., self-driving cars, automatic trading or health diagnostics this seems even more curious because it means such complex questions are addressed reductionistically dissecting the original problem into smaller subproblems rather than addressing them as a whole. Metaphorically, this may be considered as maturing process of AI settling after a rebellious adolescence against the limitations of existing fields like control theory, signal processing or statistics (Russell and Norvig, 2016). Whether it will remain this way, remains to be seen in the future.

Finally, if one considers also novel extensions for all base methods from AI, ML and statistics one can summarize the current state of these fields as follows:

Current domain of artificial intelligence, machine learning and statistics \Rightarrow Simple questions for high-dimensional data

This means that all fields seem to converge to simple question for high-dimensional data.

5. CONCLUSIONS

In this paper, we discussed the desired and current status of AI and clarified its goals. Furthermore, we put AI into perspective alongside machine learning and statistics and identified similarities and differences. The most important results can be summarized as follows:

- (1) currently, no generally accepted definition of “intelligence” is available. \Rightarrow AI remains mathematically undefined, almost 65 years after its formal inception.
- (2) the aspirations of AI are very high focusing on ambitious goals. \Rightarrow AI is not explicitly data focused—in contrast to data science.
- (3) general AI methods do not provide neurobiological models of brain functions. \Rightarrow AI methods are merely means to analyze data - similar to methods from machine learning and statistics.
- (4) addition: Also deep neural networks also do not provide neurobiological models of brain functions. \Rightarrow They are merely means to analyze data.
- (5) the current most successful AI methods, i.e., deep neural networks, focus on simple questions (classification, regression) and high-dimensional data. \Rightarrow This goes counter traditional AI but is similar to contemporary machine learning and statistics.
- (6) AI methods are not derived from a common mathematical formalism but have been introduced separately and independently. \Rightarrow There is no common conceptual framework that would unite the ideas behind different AI methods.

Finally, we would like to note that the closeness to applications of AI is certainly good for making the field practically relevant and for achieving an impact in the real world. Interestingly, this is very similar to a commercial product. A downside is that AI also comes with slogans and straplines used for marketing reasons just as those used for regular commercial products. We hope our article can help people look beyond the marketing definition of AI to see what the field is actually about from a scientific perspective.

AUTHOR CONTRIBUTIONS

FE-S conceived the study. All authors contributed to the writing of the manuscript and approved the final version.

FUNDING

MD thanks the Austrian Science Funds for supporting this work (project P30031).

REFERENCES

- Bahrammirzaee, A. (2010). A comparative survey of artificial intelligence applications in finance: artificial neural networks. *Expert System and Hybrid Intelligent Systems. Neural Comput. Applic.* 19 (8), 1165–1195. doi:10.1007/s00521-010-0362-z
- Block, N. (1981). Psychologism and behaviorism. *Phil. Rev.* 90 (1), 5–43. doi:10.2307/2184371
- Breiman, L. (2001). Random forests. *Mach. Learn.* 45, 5–32. doi:10.1023/A:1010933404324
- Brooks, R. (1991). New approaches to robotics. *Science* 253, 1227–3210. doi:1126/science.253.5025.1227
- Chen, T., and Guestrin, C. (2016). “Xgboost: a scalable tree boosting system,” in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, Seattle, Washington, University of Washington, 785–794. doi:10.1145/2939672.2939785
- Cox, D. R. (1972). Regression models and life-tables. *J. Roy. Stat. Soc. B* 34 (2), 187–202. doi:10.1007/978-1-4612-4380-9_37
- Crevier, D. (1993). *AI: the tumultuous history of the search for artificial intelligence*. New York, NY: Basic Books, 432.
- Dunn, P. K., and Smyth, G. K. (2018). *Generalized linear models with examples in R*. New York, NY: Springer, 562.
- Emmert-Streib, F., and Dehmer, M. (2019a). Defining data science by a data-driven quantification of the community. *Mach. Learn. Knowl. Extr.* 1 (1), 235–251. doi:10.3390/make1010015
- Emmert-Streib, F., and Dehmer, M. (2019b). High-dimensional lasso-based computational regression models: regularization, shrinkage, and selection. *Mach. Learn. Knowl. Extr.* 1 (1), 359–383. doi:10.3390/make1010021
- Emmert-Streib, F., and Dehmer, M. (2019c). Large-scale simultaneous inference with hypothesis testing: multiple testing procedures in practice. *Mach. Learn. Knowl. Extr.* 1 (2), 653–683. doi:10.3390/make1020039
- Emmert-Streib, F., and Dehmer, M. (2019d). Understanding statistical hypothesis testing: the logic of statistical inference. *Mach. Learn. Knowl. Extr.* 1 (3), 945–961. doi:10.3390/make1030054
- Emmert-Streib, F., Yang, Z., Feng, H., Tripathi, S., and Dehmer, M. (2020). An introductory review of deep learning for prediction models with big data. *Front. Artif. Intell.* 3, 4.
- Emmert-Streib, F. (2003). “Aktive computation in offenen systemen,” in *Lerndynamiken in biologischen systemen: vom netzwerk zum organismus*. Ph.D. thesis. Bremen, (Germany): University of Bremen.
- Emmert-Streib, F. (2006). A heterosynaptic learning rule for neural networks. *Int. J. Mod. Phys. C* 17 (10), 1501–1520. doi:10.1142/S0129183106009916
- Fan, J., Fang, L., Wu, J., Guo, Y., and Dai, Q. (2020). From brain science to artificial intelligence. *Engineering* 6 (3), 248–252. doi:10.1016/j.eng.2019.11.012
- Farcomeni, A. (2008). A review of modern multiple hypothesis testing, with particular attention to the false discovery proportion. *Stat. Methods Med. Res.* 17 (4), 347–388. doi:10.1177/0962280206079046
- Feigenbaum, E. A. (1963). Artificial intelligence research. *IEEE Trans. Inf. Theor.* 9 (4), 248–253.
- Freund, Y., and Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.* 55 (1), 119–139. doi:10.1007/3-540-59119-2_166
- Frieden, B. R., and Frieden, R. (1998). *Physics from fisher information: a unification*. Cambridge, England: Cambridge University Press, 318.
- Goertzel, B., and Pennachin, C. (2007). *Artificial general intelligence*. New York, NY: Springer, 509.

- Goldstein, H., Poole, C., and Safko, J. (2013). *Classical mechanics*. London, United Kingdom: Pearson, 660.
- Hayes-Roth, F., Waterman, D. A., and Lenat, D. B. (1983). *Building expert system*. Boston, MA: Addison-Wesley Longman, 119.
- Hebb, D. (1949). *The organization of behavior*. New York, NY: Wiley, 335.
- Hinton, G. E., Osindero, S., and Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural Comput.* 18 (7), 1527–5410. doi:1162/neco.2006.18.7.1527
- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9 (8), 1735–8010. doi:1162/neco.1997.9.8.1735
- Hoehndorf, R., and Queralt-Rosinach, N. (2017). Data science and symbolic ai: synergies, challenges and opportunities. *Data Sci.* 1 (1–2), 27–38. doi:10.3233/DS-170004
- Hodgkin, A., and Huxley, A. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* 117, 500–4410. doi:1113/jphysiol.1952.sp004764
- Hoi, S. C., Sahoo, D., Lu, J., and Zhao, P. (2018). Online learning: a comprehensive survey. Preprint repository name [Preprint]. Available at: arXiv:1802.02871 (Accessed February 8, 2018).
- Kaelbling, L., Littman, M., and Moore, A. (1996). Reinforcement learning: a survey. *J. Artif. Intell. Res.* 237–285.
- Kaplan, A., and Haenlein, M. (2019). Siri, Siri, in my hand: who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Bus. Horiz.* 62 (1), 15–25. doi:10.1016/j.bushor.2018.08.004
- Kleinbaum, D., and Klein, M. (2005). *Survival analysis: a self-learning text, statistics for biology and health*. New York, NY: Springer, 590.
- Kleinbaum, D. G., Dietz, K., Gail, M., Klein, M., and Klein, M. (2002). *Logistic regression*. New York, NY: Springer, 514.
- Koenig, S., and Simmons, R. G. (1994). "Principles of knowledge representation and reasoning," in Proceedings of the fourth international conference (KR '94), June 1, 1994, (Morgan Kaufmann Publishers), 363–373. doi:10.5555/3087304
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Adv. Neural Inform. Process. Syst.* 25 (2), 1097–1105. doi:10.1145/3065386
- Kurzweil, R., Richter, R., Kurzweil, R., and Schneider, M. L. (1990). *The age of intelligent machines*. Cambridge, MA: MIT press, Vol. 579, 580.
- Kurzweil, R. (2005). *The singularity is near: when humans transcend biology*. Westminister, England: Penguin, 672.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–4410. doi:1038/nature14539
- Legg, S., and Hutter, M. (2007). Universal intelligence: a definition of machine intelligence. *Minds Mach.* 17 (4), 391–444. doi:10.1007/s11023-007-9079-x
- McCarthy, J., Minsky, M. L., Rochester, N., and Shannon, C. E. (2006). A proposal for the Dartmouth summer research project on artificial intelligence, August 31, 1955. *AI Magazine*. 27 (4), 12. doi:10.1609/aimag.v27i4.1904
- Minsky, M., and Papert, S. (1969). *Perceptrons*. New York, NY: MIT Press, 258. doi:10.5555/1095587
- Minsky, M. L. (1967). *Computation*. Prentice-Hall Englewood Cliffs, 317.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., et al. (2013). Playing atari with deep reinforcement learning. Preprint repository name [Preprint]. Available at: arXiv:1312.5602 (Accessed December 19, 2013).
- Nagumo, J., Arimoto, S., and Yoshizawa, S. (1962). An active pulse transmission line simulating nerve axon. *Proc. IRE.* 50, 2061–2071. doi:10.1109/JRPROC.1962.288235
- Nelder, J. A., and Wedderburn, R. W. (1972). Generalized linear models. *J. Roy. Stat. Soc.* 135 (3), 370–384. doi:10.2307/2344614
- Newel, A., and Simon, H. A. (1976). Completer science as empirical inquiry: symbols and search. *Commun. ACM.* 19 (3), 113–126. doi:10.1145/360018.360022
- O'Leary, D. E. (2013). Artificial intelligence and big data. *IEEE Intell. Syst.* 28 (2), 96–99. doi:10.1109/MIS.2013.39
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. Morgan-Kaufmann, 576.
- Pomeroy, J.-C. (1997). Artificial intelligence and human decision making. *Eur. J. Oper. Res.* 99 (1), 3–25. doi:10.1016/S0377-2217(96)00378-5
- Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE.* 77 (2), 257–286. doi:10.1109/5.18626
- Rosenblatt, F. (1957). *The perceptron, a perceiving and recognizing automaton project para*. Buffalo, NY: Cornell Aeronautical Laboratory.
- Roweis, S. T., and Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science* 290 (5500), 2323–2610. doi:1126/science.290.5500.2323
- Rumelhart, D., Hinton, G., and Williams, R. (1986). Learning representations by back-propagating errors. *Nature* 323, 533–536. doi:10.1038/323533a0
- Russell, S. J., and Norvig, P. (2016). *Artificial intelligence: a modern approach*. Harlow, England: Pearson, 1136.
- Schmidhuber, J. (2015). Deep learning in neural networks: an overview. *Neural Netw.* 61, 85–117. doi:10.1016/j.neunet.2014.09.003
- Schölkopf, B., and Smola, A. (2002). *Learning with kernels: support vector machines, regularization, optimization and beyond*. Cambridge, MA: The MIT Press, 644.
- Scutari, M. (2010). Learning bayesian networks with the bnlearn r package. *J. Stat. Software.* 35 (3), 1–22. doi:10.18637/jss.v035.i03
- Searle, J. R. (2008). *Mind, language and society: philosophy in the real world*. New York, NY: Basic Books, 196.
- Sheskin, D. J. (2004). *Handbook of parametric and nonparametric statistical procedures*. 3rd Edn, Boca Raton, FL: RC Press, 1193.
- Simon, H. A. (1965). *The shape of automation for men and management*. New York, NY: Harper & Row, 13, 211–212.
- Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT University Press, 123.
- Sutton, R., and Barto, A. (1998). *Reinforcement learning*. Cambridge, MA: MIT Press, 344.
- Thrun, S. (2002). *Robotic mapping: a survey, Exploring artificial intelligence in the new millennium*, 1–35. doi:10.5555/779343.779345
- Turing, A. (1950). Computing machinery and intelligence. *Mind* 59, 433–460. doi:10.1093/mind/LIX.236.433
- Vapnik, V. N. (1995). *The nature of statistical learning theory*. New York, NY: Springer, 188.
- Wang, P. (2006). *Rigid flexibility: the logic of intelligence*. New York, NY: Springer Science & Business Media, Vol. 34, 412.
- Wang, P. (2019). On defining artificial intelligence. *J. Artif. Gen. Intell.* 10 (2), 1–37. doi:10.2478/jagi-2019-0002
- Weisberg, S. (2005). *Applied linear regression*. New York, NY: John Wiley & Sons, 528, 352.
- Winston, P. H., and Brown, R. H. (1984). *Artificial intelligence, an MIT perspective*. Cambridge, MA: MIT Press, 492.
- Wolpert, D. M., and Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Netw.* 11 (7–8), 1317–1329. doi:10.1016/S0893-6080(98)00066-5
- Wooldridge, M. J., and Jennings, N. R. (1995). Intelligent agents: theory and practice. *Knowl. Eng. Rev.* 10 (2), 115–152.
- Yampolskiy, R. V., and Fox, J. (2012). "Artificial general intelligence and the human mental model," in *Singularity hypotheses*. New York, NY: Springer, 129–145.
- Yu, K. H., Beam, A. L., and Kohane, I. S. (2018). Artificial intelligence in healthcare. *Nat. Biomed. Eng.* 2, (10), 719–731. doi:10.1038/s41551-018-0305-z

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Emmert-Streib, Yli-Harja and Dehmer. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.