

Review

Machine Learning in Artificial Intelligence: Towards a Common Understanding

Niklas Kühl; Marc Goutier; Robin Hirt; Gerhard Satzger*

Karlsruhe Institute of Technology, Karlsruhe, Germany

*Corresponding author

Niklas Kühl

Karlsruhe Institute of Technology, Karlsruhe, Germany

Article information

Received: March 30th, 2024; Revised: June 22nd, 2024; Accepted: July 20th, 2024; Published: August 6th, 2024

Cite this article

Kühl N. Machine learning in artificial intelligence: Towards a common understanding. 2024; 2(2). doi: <https://doi.org/10.70705/ppp.ltcs.2024.v02.i02.pp53-59>

ABSTRACT

“Machine learning” and “artificial intelligence” have emerged as prominent applications in the last ten years. In the scientific and media communities, you could hear both phrases used interchangeably or with distinct meanings. Our goal in writing this is to define each of these words and their connection to AI, with a focus on machine learning’s role in the field. Our conceptual framework, based on a survey of the literature, elucidates machine learning’s function in the development of (artificial) intelligent agents. Therefore, our goal is to provide a foundation for future research and (interdisciplinary) conversations by clarifying key terms.

INTRODUCTION

Mark Zuckerberg emphasized the need of Facebook’s “AI tools (... to) identify hate speech (...)” or “ (...) terrorist propaganda” at his April 2018 testimony before the US senate [1]. Most of the time, researchers in the area of (supervised) machine learning would refer to the process of finding certain occurrences on social media platforms as a classification problem [2]-[4]. In spite of this, the phrase artificial intelligence (AI) is frequently used synonymously with machine learning due to AI’s growing popularity [5]. This is true not only in the aforementioned interview with Facebook’s CEO but also in many theoretical and application-oriented contributions in recent literature [7]-[9]. Even though he knows it is incorrect, Carner(2017) says he still uses AI to mean machine learning [10]. However, when discussing methodologies, ideas, and outcomes, such fuzziness could cause many mistakes in both theory and practice.

The lack of useful scientific definition, despite the terminology’s widespread usage, is remarkable. Therefore, the purpose of this study is to explain the connection between AI and machine learning. Our discussion of the function of machine learning used by AI instances, more specifically by intelligent agents. The capabilities and deployment of intelligent agents are seen via a machine learning lens in order to achieve this.

Our article has three main contributions. To begin, we extend the theoretical framework of Russel & Norvig (2015) [11] by decomposing the “thinking” layer of an intelligent agent into two

distinct sublayers: “learning” and “executing.”. The second part of our work is demonstrating how this difference allows us to separate the various AI-related machine learning contributions. A continuum between human engagement and agent autonomy is defined, thirdly, by drawing on the implementations of the execution and learning sublayers (“backend”).

The next section of this article begins with a literature survey covering important works in AI and machine learning. Our conceptual framework, which emphasizes machine learning’s contribution to AI, is then presented and elaborated. We finish with a summary, present constraints, and a forecast, and use that information to generate a research plan for the future.

2. Related work

As a base for our conceptual work, we first review the different notions, concepts, or definitions of machine learning and artificial intelligence within extant research. In addition, we elaborate in greater detail on the theories which we draw upon in our framework.

2.1. Terminology

Machine learning and artificial intelligence, as well as the terms data mining, deep learning and statistical learning are related, often present in the same context and sometimes used interchangeably. While the terms are common in different communities, their particular usage and meaning varies widely.

For example, statistical learning—a collection of methods and techniques for acquiring information, making predictions, and making choices via the construction of models from data sets—is the main emphasis in the discipline of statistics [12]. When seen through the lens of statistics, machine learning becomes an application of statistical learning [13].

One subfield of computer science known as machine learning is concerned with finding effective ways to use available computing power to address specific issues [14]. Although machine learning draws on statistical techniques, it also incorporates approaches that aren't strictly based on statisticians' work. This has led to the development of novel methods and approaches that have been well recognized and appreciated in the area [15], [16]. In particular, deep learning has gained a lot of attention recently [17]. The multi-processing layers that make up deep learning models allow them to learn data representations with varying degrees of abstraction. For example, in voice recognition [18] and picture recognition [19], deep learning has greatly enhanced machine learning's capabilities.

Data mining stands in contrast to the aforementioned concepts; it is the practice of using quantitative analytical approaches to address practical issues, often in commercial contexts [20]. Data mining, as it pertains to machine learning, is the process by which useful machine learning models are produced. The goal is not to learn more about ML algorithms; rather, it is to use them on data and draw conclusions. Data mining therefore has machine learning as its cornerstone [21]. Artificial intelligence, on the other hand, makes use of methods like statistical learning, machine learning, and using other methods, such as descriptive statistics, to imitate artificial intelligence.

The rest of this work is built upon Figure 1 and the words described in this paragraph. Nevertheless, there is contentious debate over the general language and linkages of the notions [22]. Thus, this paper's primary objective is to shed light on the jargon and, more specifically, to define machine learning's function inside artificial intelligence. We take a closer look at machine learning and AI to help you grasp the concepts better.

2.2. Machine learning

The term “machine learning” refers to a collection of methods used to train computers to solve problems automatically, rather than requiring human intervention [23]. Unsupervised and supervised machine learning are generally distinct from one another. Since the most popular approaches are supervised, we will be concentrating on the latter in this study [24]. Learning in supervised machine learning is constructing knowledge about a task by analyzing a set of instances, or “past experience” [25]. Although statistical approaches are used throughout the learning process, problem-solving does not need human modification or programming of rules or tactics. Building a model using an algorithm on a collection of known data points to acquire insight on an unknown set of data is the ultimate goal of all (supervised) machine learning approaches [11], [26].

Model initiation, performance estimation, and deployment are the three primary stages of a machine learning model's process, but the exact definitions of these phases could vary somewhat among processes [27]: A human user selects an appropriate machine learning method, describes the issue, and prepares and analyzes the data set

during the model start phase. Afterwards, a well-performing configuration is chosen based on its performance in accomplishing a certain job after the performance estimation validates numerous parameter permutations characterizing the method. As a last step, the model is implemented and used to solve the problem using data that has not been previously seen.

One important part of human cognition is learning, which is defined as “the processes by which the sensory input is transformed, reduced, elaborated, stored, recovered, and used” [28, p. 4]. In order to make sense of incoming data, humans use abstract knowledge, which allows them to analyze massive amounts of data. Machine learning models may, in a limited way, imitate human intelligence because of their adaptability.

On the other hand, machine learning is just a collection of techniques that may be used to create analytical models that can be integrated into bigger IT artifacts by learning patterns in existing data.

2.3. Artificial intelligence

A number of academic fields have contributed to our understanding of artificial intelligence (AI), including computer science [18,19], philosophy [20,21], and futures studies [22, 23]. Since computer science is the most applicable discipline for distinguishing between AI and machine learning and for determining the role of machine learning in the former, it is the primary emphasis of this study.

There are many subfields within artificial intelligence research [11]. They diverge, first, in their focus on the goal of AI applications (thinking vs. acting), and second, in the nature of the decisions they want to make (pursuing a human-like choice vs. an ideal, logical conclusion). As seen in Table 1, this differentiation gives rise to four separate lines of inquiry.

Artificial intelligence (AI) must be a computer with a mind, according to the “Cognitive Modeling” (thinking humanly) stream [34]. Furthermore, this entails carrying out human-level reasoning [35], using the same inputs as a person would and arriving at the same conclusion [36] based on the same reasoning stages.

According to the “Laws of Thought” school of thought, which emphasizes logical reasoning, AI must always choose the best course of action, regardless of how a human might respond.

Computational models [37] that represent logic are therefore necessary for AI to adhere to the rules of cognition.

That artificial intelligence (AI) must behave intelligently while interacting with people is implied by the “Turing Test” (i.e. behaving humanly) stream. Artificial intelligence (AI) has to be able to do human-level work to achieve these goals [38]. The Turing Test can verify these conditions [39].

The “Rational Agent” school of thought concludes by viewing AI through the lens of an intelligent [40] or rational [11] agent. Not only does this agent behave independently, but it does so with the goal of attaining the logically optimal result.

Determining what intelligence is and then applying that definition to the development of intelligent robots is another possible approach to defining AI. A measurement of intelligence is defined by Legg

and Hutter [41] using psychological criteria, theories of human intellect, and intelligence tests. Their formulation is the basis of an agent-environment paradigm that they use to characterize intelligence and, more specifically, artificial intelligence when the agent in question is a computer. The “acting rationally” stream is quite close to their concept.

Research in the area of artificial intelligence also focuses on how to categorize AI systems, in addition to trying to define AI in general. The Searle

Separating weak AI from strong AI is recommended in [42]. A powerful AI really has mental states, unlike a weak AI that only acts like it thinks. On the other hand, Gubrud [43] classifies AI according to the nature of the job. In each given subject, an artificial general intelligence (AGI) system can perform at least as well as a human brain, but it doesn’t need any specific training or instruction from a person.

state of mind. A narrow AI, on the other hand, is one that compares well to or even surpasses the human brain in very restricted domains [44].

Because of its significance in relation to AI machine learning implementation, the “Rational Agent” stream will be investigated more below. In Section 3, we will revisit the other three lines of inquiry and demonstrate their compatibility with our agent-based AI system.

The “Rational Agent” school of thought holds that agent behavior is the ultimate manifestation of intelligence. These beings exhibit five traits: they “operate autonomously, perceive their environment, persist over a prolonged time period, adapt to change, and create and pursue goals” [11, p. 4]. Instead of acting in isolation, an agent’s actions are defined in relation to the environment in which it operates. The robot is able to identify its surroundings by its sensors, process the data provided by its agent software, and then carry out an action using its actuators. Acting in a way that maximizes the predicted result according to this performance metric—based on the agent’s present and historical knowledge of the environment and the alternative actions—is another need for being a rational agent.

According to Russel and Norvig, there are four distinct kinds of agents in the agent program [11] when it comes to the broad definition of agents: The difference between a basic reflex agent and a model-based reflex agent is that the latter takes the agent’s internal state into account in addition to sensor input when making a reaction. In order to accomplish its objectives, a goal-based agent chooses the optimal course of action. An objective may be satisfied in one of two ways: either by achieving it or by failing to do so. A utility-based agent, in contrast, seeks to maximize a full utility function rather than a discrete one. A program may be extended to make an agent a learning agent. In this kind of learning agent, one part is responsible for making decisions based on sensor data, while the other part learns from its surroundings, creates its own issues, and, if feasible, makes improvements to the performance part.

An agent, their surroundings, and an objective make up the agent-en-

vironment paradigm. A person’s intelligence may be defined as their “capacity to achieve goals in a wide range of environments” [41, p. 12]. Perceptions created by the environment provide input to the agent. Some perceptions are just observations of the world around the agent, while others are reward signals that show how well the agent has done in achieving its objectives. The agent takes these signals as input, processes them, and then relays the results to its surroundings.

3. A framework for understanding the role of machine learning in artificial intelligence

In order to understand the interplay of machine learning and AI, we base our concept on the framework of Russel & Norvig [11]. With their differentiation between the two objectives of AI application, acting and thinking, they lay an important foundation.

3.1. Layers of agents

Machine learning’s place in artificial intelligence may be better grasped from a vantage point that focuses on the creation and deployment of smart agents. This viewpoint is essential because it helps us to relate the various parts of machine learning to the capacities of smart agents. Looking at the thinking and acting capabilities of an intelligent agent and applying this to software architecture, we can say that the acting capabilities are the frontend and the thinking portion is the backend. For the sake of autonomy, flexibility, and parallel development, software developers normally keep form and function completely distinct [45]. The environment’s interface is known as the frontend. Many forms are possible. A humanoid template with expanded expressive capabilities [48], a machine-readable web interface [46], or a human-readable application [47] are all possible forms of intelligent agent interfaces. Technically speaking, the frontend can’t communicate with its surroundings without sensors and actuators. Through the frontend, sensors report any changes or events in the environment, which are then sent to the backend. For example, they may decipher the temperature readings from a manufacturing equipment [49] or interpret the graphics from a human interaction [50]. In contrast, components that move and regulate a mechanism are known as actuators. In contrast to sensors, which only collect data, actuators do something, such automatically purchasing stocks [51] or modifying a humanoid’s facial expressions [52]. To put an agent’s artificial intelligence to the test of how well it mimics human behavior, one may say that the Turing test [39] occurs when the environment interacts with the frontend, specifically when sensors and actuators work together. Although every frontend has sensors and actuators, the specific frontend design is irrelevant to our work. What is crucial is the existence of a backend-independent, encapsulated frontend.

The backend provides the necessary functionalities, which depict the thinking capabilities of an intelligent agent. Therefore, the agent needs to learn and apply learned knowledge.

In consequence, machine learning is relevant in this implementation layer. When regarding the case of supervised machine learning, we need to further differentiate between the process task that is building (=training) adequate machine learning models [21] and the process task that is executing the deployed models [53]. Therefore, to further understand the role of machine learning within intelligent agents, we refine the thinking layer of agents into a learning sublayer (model building) as well as an executing sublayer (model execution)². Hence, we regard the necessary implementation for the learning sublayer as

the learning backend, while the executing sublayer is denoted by the executing backend.

3.2. Types of learning

To begin with, the learning backend determines whether or not the intelligent agent can learn. Then, it specifies how the agent may learn, including the exact algorithms it employs, the data processing type it employs, and the notion

is dealt with, etc., drift [54]... Thus, we use the language of Russel and Norvig [11] and distinguish between two categories of intelligent agents: simple-reflex agents and learning agents. This distinction is particularly relevant when looking at AI from a machine learning standpoint, as it determines whether the thinking layer's underlying models are static (simple-reflex) or dynamic (learning). Both may be found in more modern works of literature. Using a hospital pneumonia early warning system as an example of a simple-reflex agent, Oroszi and Ruhland [55] construct and test a plausible agent model, although the adaptive learning that occurs after deployment may be crucial. Anaphora resolves [56], pedestrian prediction, and other domains also often use agents with single-trained models.

or annotating objects [58]. Additionally, instances for learning agents may be found in more recent research. With an emphasis on constantly constructing and updating models inside agents, Mitchell et al. [59] introduce the idea of “never-ending learning” agents. Liebman et al. (2019) construct a self-learning bot designed to suggestions for musical playlists [60]. Some more examples include controlling heat pump thermostats [61], understanding the meaning of words [63], or an agent acquiring collective knowledge across several jobs [62].

Both the agent's general design and the impact of machine learning are affected by the decision on this feature in general (simple-reflex vs. learning agent). The second graphic shows the overall structure that we were able to construct. Finally, in the execution sublayer of a simple-reflex agent, machine learning occurs as a once-trained model. On the other hand, it helps a learning agent's model in the execution sublayer to become better over time. The execution layer is responsible for acquiring environmental information and feedback, which is the basis of this improvement.

3.3. Continuum between human involvement and machine involvement

When it comes to the executing backend and the learning backend, it is not only of importance if and how underlying machine learning models are updated—but how much automated the necessary processes are. Every machine learning task involves various process steps, including data source selection, data collection, preprocessing, model building, evaluating, deploying, executing and improving (e.g. [21], [53], [64]). While a discussion of the individual steps is beyond the scope of this paper, the autonomy and the automation of these tasks as an implementation within the agent is of particular interest in each necessary task of the machine learning lifecycle [27].

For instance, while the execution of a once-built model can be fairly easily automated, the automated identification of an adequate data source for a new problem or retraining as well as a self-induced model building are more difficult. Therefore, we need to view the human involvement in the necessary machine learning tasks of an intelligent agent, as depicted in figure 3. While it is hard to draw a clear line between all possible forms of human involvement in the machine learning-relevant tasks of an intelligent agent, we see this phenomenon rather as a continuum. The continuum ranges between none or little agent autonomy with full human involvement (e.g. [65]– [67]) on the one extreme as well as the full agent

autonomy and no or little human involvement for the delivered task on the other (e.g. [68]–[70]). For example, an intelligent agent with the task to autonomously drive a car considering the traffic signs already proves a high degree of agent autonomy. However, if the agent is confronted with a new traffic sign, the learning of this new circumstance might still need human involvement as the agent might not be able to “completely learn by itself” [71]. Therefore, the necessary involvement of humans, especially in the thinking layer (= executing backend and learning backend), is of major interest when describing AI and the underlying machine learning models. The degree of autonomy for each step of machine learning can be investigated and may help to characterize the autonomy of an agent in terms of the related machine learning tasks.

4. Research priorities for machine-learning-enabled artificial intelligence

Still in its conceptual stages is the offered framework of machine learning and its function inside intelligent agents. Further study might clarify the language and explore unknown area for machine-learning enabled artificial intelligence, given the misconceptions and ambiguity of the two phrases [6-9].

The first step is to construct the framework iteratively and empirically. In order to test the efficacy of the framework, we must first locate several examples of intelligent beings in diverse fields. It would be fascinating to see the mapping between the framework and AI initiatives that use machine learning, both in academia and in the real world. It would be even more interesting to quantify the percentage of these projects that use learning agents compared to those that do not. In addition, these kinds of examples would be useful for learning about the human element required by modern intelligent agents, which would allow us to ascertain the “degree” of autonomy these agents need in relation to all parts of their operation (learning, acting, and executing).

Secondly, one area that should be improved is the need to engage less people. Our previous discussion has shown that we see this spectrum as a continuous line between human intervention and agent autonomy. There are two options that spring to mind right away. The concepts and techniques of transfer machine learning center on the idea of moving models and other forms of information from one setting to another [72]. As further studies in this area uncover potential and application-oriented methods to use transfer, this might definitely aid in reducing human participation.

use machine learning to automatically adjust to new or altered duties [73].

Not only is it important to consider the initial model construction process, but also how to handle changes in the environment, when

thinking about models that have already been deployed as part of the backend-layer. There are a lot of ideas in the so-called “concept drift” sector on how to adjust models and identify changes, but there haven’t been many successful applications [54], [74].

5. Conclusion

The purpose of this study is to define machine learning and its function in AI, with a focus on intelligent agents. We provide a paradigm that emphasizes the two examples of learning agents and simple-reflex agents, and the potential role of machine learning in both. A basic reflex agent is an intelligent agent that uses machine learning models as pre-trained models and is unable to learn anything new from its surroundings. We refer to this layer of executing knowledge as the executing backend when discussing implementation details. In this scenario, the agent may make use of machine learning models that have already been created, but it cannot create or update its own models. A learning agent, on the other hand, is one that can take in data from its surroundings and use that information to refine the machine learning models used by the execution sublayer. In order for learning agents to make use of machine learning for model creation and training, they have an extra sublayer called the learning backend.

The level of independence needed by the agent’s machine learning must be captured before these two sublayers may be implemented. The human element is the main emphasis here, as it pertains to the data collecting and algorithm selection processes that are fundamental to machine learning.

The current study has various limitations and is currently in its conceptual stage. To begin, while the suggested framework does help us better comprehend machine learning in AI, we still need further empirical research to determine how well current AI applications that make use of machine learning fit into this scheme. To verify the model and assess its amount of depth, AI designers might be interviewed by experts. In addition, to better comprehend the level of autonomy of state-of-the-art agents, we need to discover techniques to measure the human engagement in AI activities connected to machine learning.

Our paradigm is still in its infancy, but it should help researchers and practitioners use more accurate language when talking about AI and machine learning. It stresses how crucial it is to avoid using the phrases while simultaneously highlighting the function of machine learning in the context of a certain agent implementation.

REFERENCES

In 2018, “The Washington Post” published a transcript of Mark Zuckerberg’s Senate testimony. [On the web]. Mark Zuckerberg’s Senate Hearing Transcript: https://www.washingtonpost.com/news/the-switch/wp/2018/04/10/transcript-of-mark-zuckerbergs-senate-hearing/?utm_term=.4720e7f10b41. [Monthly updated:15-Jun-2018].

To wit: “Hateful Symbols or Hateful People?” by Z. Waseem and

D. Hovy. Proceedings of the 2016 NAACL Student Research Workshop, pages 88–93,

presents predictive features for Twitter hate speech detection. the third [Detecting hate speech on the world wide web] (W. Warner and J. Hirschberg, Proceeding LSM ‘12 Proc. Second Work). Journal of Language and Social Media, volume LSM, 2012, pages 19–26.

[4] “Uncovering the dark web a case study of jihad on the web,” International Review of Research on Open and Distance Learning, vol. 14, no. 4, pp. 90-103, 2013, written by H. Chen, W. Chung, J. Qin, E. Reid, M. Sageman, and G. Weimann.

Economic study and Policy, vol. 58, no. 5, “Trends and priority shifts in artificial intelligence technology invention: A global patent analysis” (H. Fujii and S. Managi, 2015). page 60–69, year 2018.

The interview transcript from 2016 with Mark Zuckerberg, “How to Build the Future,” was published by the University of Wisconsin.

Data Protection Act and General Data Protection Regulation, 2017, “Big Data, artificial intelligence, machine learning and data protection” (in the Information Commissioner’s Office’s working title). [On the web]. To access the paper, go to: <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>. (June 15, 2018) as accessed.

[8] “Big Data Management, Access, and Protection,” by J. A. Brink, published in 2017 in the Journal of the American College of Radiology, volume 14, issue 5, pages 579-580.

[9] Thomas Nawrocki, Peter D. Maldjian, Sarah E. Slasky, and Susan E. Has the Dystopia of the Radiologist Been Dramatically Exaggerated?—Artificial Intelligence and Radiology, G. Contractor, Academic Radiology, 2018.

[10] “Artificial intelligence and behavioral economics” by C. F. Camerer, Economics of Artificial Intelligence, University of Chicago Press, 2017.

[11] AI: A Modern Approach, 3rd ed., 2015, S. J. Russell and P. Norvig.

[12] “The elements of statistical learning: data mining, inference and prediction,” Math. Intell., vol. 27, no. 2, pp. 83-85, 2005, by T. Hastie, R. Tibshirani, J. Friedman, and J. Franklin.

The authors of this work are Bousquet, von Luxburg, and Rätsch. It is titled “Advanced Lectures on Machine Learning: ML Summer Schools 2003” and was held in Tübingen, Germany from February 2–14, 2003 in Canberra, Australia.

Volume 3, Section 3, pages 4-16, 2003. The publisher, Springer, 2011.

In their work, Mohri, Rostamizadeh, and Talwalkar [14] demonstrate Foundations of machine learning. 2012 by MIT Press.

“Extreme learning machine: a new learning scheme of feed-forward neural networks,” in Neural Networks, 2004, by G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew. Proceedings. in 2004 The 2018 IEEE International Joint Volume 2, pages 985–990, 2004. Conference on.

The paper “Machine learning in automated text categorization” by F. Sebastiani was published in the ACM Computing Surveys in 1991. 2002, pages 1–47.

In their 2015 article “Deep learning,” Y. A. LeCun, Y. Bengio, and G. E. Hinton discuss deep learning.

Referenced in “Deep Neural Networks for Acoustic Modeling in Speech Recognition” published in 2012 by IEEE Signal Processing Magazine, the authors include G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury.

“Deep Residual Learning for Image Recognition,” presented at the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), was written by K. He, X. Zhang, S. Ren, and J. Sun.

20. C. Schommer, “An Unified Definition of Data Mining,” CoRR, vol. abs/0809.2696, 2008.

[21] Data Mining: Practical Machine Learning Tools and Techniques, Third Edition, vol. 54, no. 2, 2011, by I. H. Witten, E. Frank, and M. a. Hall.

The question “What distinguishes data mining, statistics, machine learning, and AI?” was cross-validated [22]. year 2014.

Automatic Topology and Size Design of Analog Electrical Circuits Using Genetic Programming, by J. R. Koza, F. H. Bennett, D. Andre, and M. A. Keane, published in Artificial Intelligence in Design ‘96, 1996.

[24] “Machine learning: Trends, perspectives, and prospects,” Science, 2015. by M. I. Jordan and T. M. Mitchell.

Machine Learning, volume 1, issue 1, 1997, by T. M. Mitchell [25].

Data mining, inference, and prediction (vol.9), edited by T. Hastie, R. Tibshirani, and J. Friedman, published by Springer in 2017, is referenced in [26].

The paper “An end-to-end process model for supervised machine learning classification: from problem to deployment in information systems” was presented at the 2017 DESRIST Research-in-Progress conference by R. Hirt, N. Kühl, and G. Satzger.

[28] In Cognitive Psychology, by U. Neisser, 1967.

[29] N. J. Nilsson, AI: A New Synthesis, 1998, vol. 125, no. 1-2.

[30] “The Logic of Action,” The Stanford Encyclopedia of Philosophy, 2016, eds. K. Segerberg, J.-J. Meyer, and M. Kracht. [On the web]. See: <https://plato.stanford.edu/archives/win2016/entries> The document was accessed on June 15, 2018.

[31] “Logic and Artificial Intelligence” by R. Thomason, 2016 edition, Stanford Encyclopedia of Philosophy. [On the web]. On the market:

Here are the entries: <https://plato.stanford.edu/archives/win2016/>. [Checked on: 15-Jun-2018] /logic-ai.

Section on Technology Pennsylvania: Penny Hill Press and the National Science and Technology Council, “Preparing for the future of Artificial Intelligence,” Committee on Technology Chapter 58, published by Penny Hill Press and the National Science and Technology Council. 2016 by CreateSpace, an independent publishing platform.

In a study conducted by P. Stone, R. Brooks, E. Brynjolfsson, R. Calo, O. Etzioni, G. Hager, J. Hirschberg, S. Kalyanakrishnan, E. Kamar, S. Kraus, K. Leyton- Brown, D. Parkes, W. Press, A. Saxenian, and J. Shah,

“Artificial Intelligence and Life in 2030,” in One Hundred Year Study Artif. Intell. Rep. 2015-2016 Study Panel, pp. 52–53, 2016, authors M. Tambe and A. Teller.

The Very Idea of Artificial Intelligence, by J. Haugeland [34]. 1989, MIT Press.

The ability of computers to think is discussed in R. Bellman’s An Introduction to Artificial Intelligence (35). Fraser and Boyd, 1978.

[36] “GPS, a program that simulates human thought,” 1961, by A. Newell and H. A. Simon.

Referenced as “Introduction to artificial intelligence” in the 1985 International Journal of Adaptive Control and Signal Processing, by D. McDermott and E. Charniak, volume 2, issue 2, pages 148–149.

“Artificial intelligence,” by E. Rich and K. Knight, [38] 1991 by

McGraw-Hill.

[39] “Computing Machine and Intelligence” by A. M. Turing, published in 1950 in MIND, volume 6, issue 236 (pp. 433–460).

[40] “Computational Intelligence and Knowledge,” published in 1998 in Computer Science: A Logical Approach, edited by D. L. Poole, A. Mackworth, and R. G. Goebel.

[41] In their 2007 article “Universal intelligence: A definition of

machine intelligence,” published in Minds Mach., S. Legg and M. Hutter explore the concept of AI and its implications.

In 1980, J. R. Searle published a book titled “Minds, Brains, and Programs” in volume 3, pages 1–19.

In his 1997 paper “Nanotechnology and international security,” M. A. Gubrud addressed this same topic at the Fifth Foresight Conference on Molecular Nanotechnology.