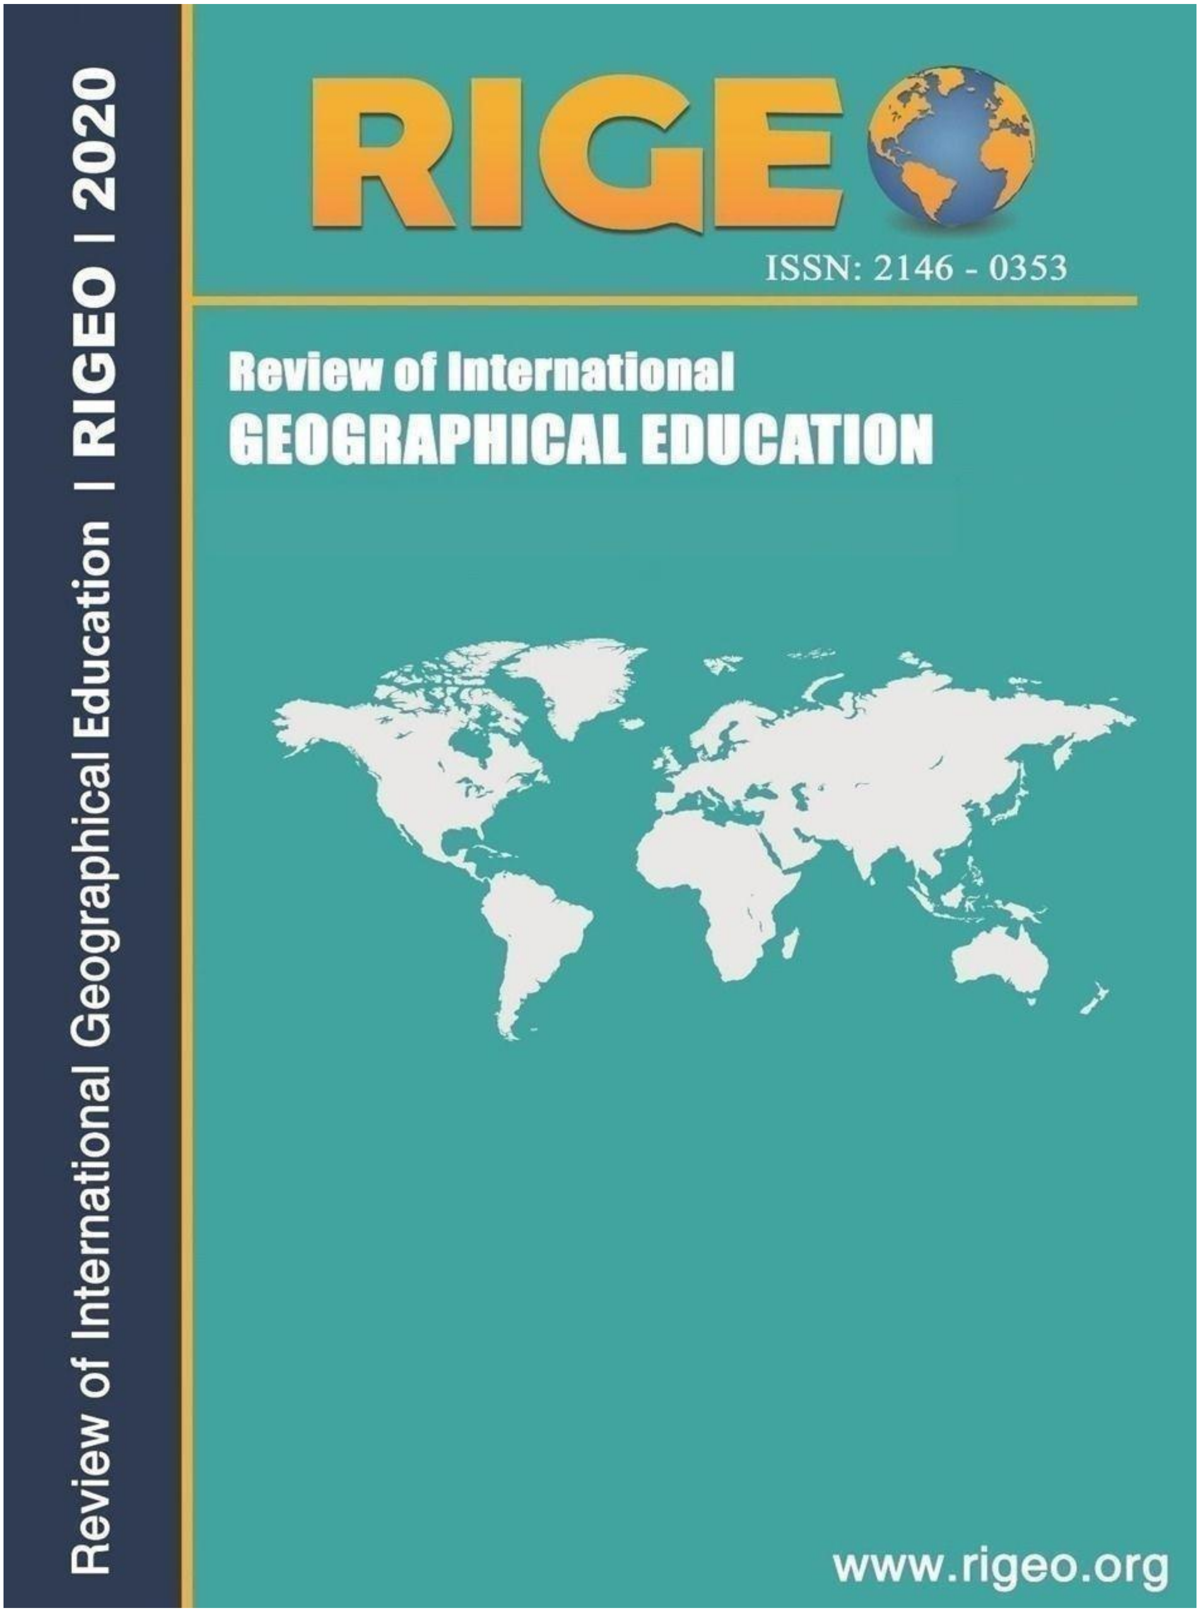
Review of International Geographical Education | RIGEO | 2020
RIGEO
ISSN: 2146 - 0353
Review of International
GEOGRAPHICAL EDUCATION
www.rigeo.org

1

Towards a Unified Framework: Machine Learning's Role in Artificial Intelligence

Dr. GANDHAVALLA RAO SAMBASIVARAO[1], Dr. MOHD ATEEQ AHMED[2]

Department of Information technology
Nawab Shah Alam Khan College of Engineering and Technology (NSAKCET)

Abstract

In the past ten years, "machine learning" and "artificial intelligence" have become popular ways to use technology. In research and the media, both terms are often used, sometimes with the same meaning and sometimes with different ones. The goal of this work is to make the connection between these terms clearer and, in particular, to explain what machine learning has to do with artificial intelligence. As part of our review of the relevant literature, we present a conceptual framework that makes the role of machine learning in creating (fake) intelligent creatures more clear. So, we want to make the terms more clear and give people from different fields a place to start talking and doing more study.

1. Introduction

In April 2018, Mark Zuckerberg spoke in front of the US Senate and talked about how important it was for Facebook's "AI tools" to be able to spot "hate speech" or "terrorist propaganda" [1]. Labelling specific posts on social media sites is something that researchers usually talk about as a classification job in the field of (supervised) machine learning [2, 4]. But because artificial intelligence (AI) is becoming more popular [5], the terms AI and machine learning are often used interchangeably. This is true not only by Facebook's CEO in the example above and in other interviews [6], but also in a lot of recent theoretical and application-focused literature [7]–[9]. Carner (2017) even says that he still uses AI to mean machine learning, even though he knows this isn't right [10]. But this kind of vagueness can lead to a lot of mistakes in study and practice when people talk about methods, ideas, and results.

It's strange that, even though the terms are used a lot, there isn't much useful science definition. The goal of this study is to explain what machine learning and artificial intelligence have to do with each other. We talk more about the part of Machine learning is used in instances of artificial intelligence, more specifically in intelligent robots. To do this, we look at the powers of intelligent agents and how they are implemented from the point of view of machine learning.

Our study makes three kinds of contributions. As a first step, we build on Russel and Norvig's (2015) [11] theory approach by breaking down the "thinking" layer of any intelligent agent into two different sublayers: "learning" and "executing." Second, we show how this division lets us tell the difference between the different ways that machine learning can help intelligent creatures. Third, we use how the processing and learning sublayers are implemented (the "backend") to set a range of levels of human participation and AI liberty.

First, we will look at important works from the fields of machine learning and artificial intelligence in the rest of this study. Next, we share and expand on our conceptual framework that shows how machine learning has helped artificial intelligence. Based on this, we come up with a plan for future study and end with an overview, a list of the problems we've found so far, and an outlook.

2. Related work

Before we start our mental work, we look at the different ideas, concepts, or meanings of artificial intelligence and machine learning that have already been studied. In addition, we go into more depth about the ideas that we use to build our system.

2.1. Terminology

This includes artificial intelligence, data mining, deep learning, statistical learning, and machine learning. These terms are often used together and in the same sentence. The terms are used in many places, but they have very different meanings and ways of being used.

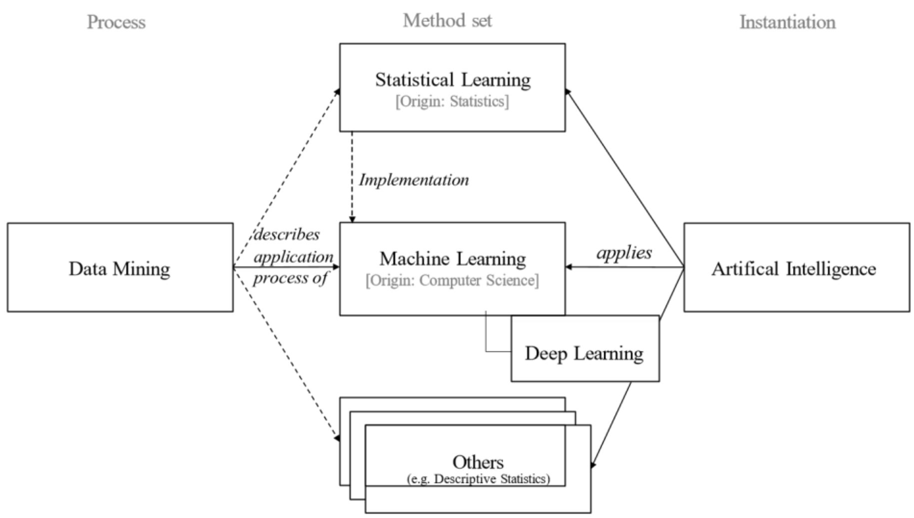

Figure 1. General terminology used in this paper

For instance, in the field of statistics the focus is on *statistical learning*, which is defined as a set of me- thods and algorithms to gain knowledge, predict outcomes, and make decisions by constructing models from a data set [12]. From a statistics point of view, machine learning can be regarded as an implemen- tation of statistical learning [13].
Within the field of computer science, *machine learning* has the focus of designing efficient algorithms to solve problems with computational resources [14]. While machine learning utilizes approaches from statistics, it also includes methods which are not entirely based on previous work of statisticians—resulting in new and well-cited contri- butions to the field [15], [16]. Especially the method of deep learning raised increased interest within the past years [17]. *Deep learning* models are composed of multiple processing layers which are capable of learning representations of data with multiple levels of abstraction. Deep learning has drastically improved the capabilities of machine learning, e.g. in speech [18] or image recognition [19].

In contrast to the other terms, data mining is the process of using quantitative analysis methods to answer problems in the real world, such as those that come up in business [20]. When it comes to machine learning, data mining is the process of making models that make sense. It's not important to learn more about machine learning techniques; what matters is using them on data to learn something. Because of this, machine learning can be seen as the basis for data mining [21]. Artificial intelligence, on the other hand, uses methods such as machine learning and statistical learning or other methods, like descriptive statistics, to make machines act like they are smart. Everything else in this paper is based on Figure 1 and the words defined in this text. But there is disagreement about the general language and connections between the ideas [22]. So, the main point of this study is to help you understand the terms better and, more specifically, to make sense of what machine learning means in AI. To get a better sense of what machine learning and AI mean, we look at both of them in more depth.

2.2. Machine learning

A group of methods called "machine learning" are used to fix many different kinds of problems in the real world by teaching computers how to solve them without being told just what to do [23]. We can tell the difference between controlled and unsupervised machine learning in most situations. Here, we'll focus on the second one because controlled methods are the ones that are most commonly used [24]. When it comes to guided machine learning, learning means using a set of examples (called "past experience") to get better at a job [25]. Statistical methods are used to help people learn, but rules or strategies to solve a problem don't have to be changed or programmed by hand. It is the goal of supervised machine learning methods to build a model by running an algorithm on a set of known data points in order to learn more about a set of unknown data [11], [26].

So, the steps used to "create" a machine learning model are a little different in how they define them, but there are usually three main ones: starting the model, estimating its performance, and deploying it [27]. A person describes a problem, prepares and processes a data set, and picks a machine learning method that will work for the job during the model start phase. Next, during the performance estimate, different parameter sets that describe the algorithm are checked to make sure they work, and the best setup is chosen based on how well it solves a certain job. Finally, the model is implemented and used to solve the problem on data that hasn't been seen yet.

Human learning is an important part of how we think and remember things [28, p. 4]. Learning includes all the steps our senses go through to change, simplify, build on, store, retrieve, and use information. People handle a lot of data by using general knowledge that helps them understand what they are being told. Because they can change over time, machine learning models can imitate the thinking skills of a human being on their own.

However, machine learning is just a group of techniques for finding trends in current data. This creates analysis models that can be used in bigger IT projects.

2.3. Artificial intelligence

Artificial intelligence (AI) comes from a lot of different areas of research, like computer science [18, 19], philosophy [20, 21], and futures studies [22, 23]. We mostly look at computer science in this work because it is the most useful for figuring out what machine learning has to do with AI and how the two terms are different.

AI study can be split up into different lines of work [11]. One difference between these lines is the goal of AI use (thinking vs. action), and the other is the type of decision making (aiming for a human-like decision vs. an ideal, reasonable decision). This difference leads to four lines of study, which can be seen in Table 1.

The "Cognitive Modelling" (thinking like a person) stream says that an AI must be a machine that has a mind [34]. It also means thinking like a human [35], based on both the same output as a human would get from the same input and the same steps of reasoning that led to the ending [36].

The "Laws of Thought" stream, which means "thinking logically," says that an AI has to make the logical choice, even if a person gives a different answer.

Table 1. AI research streams based on Russell & Norvig [11]

Objective / Application to	Humanly	Rationally
Thinking	Cognitive Modeling	“Laws of thought “
Acting	Turing Test	Rational Agent

So, for AI to work, it needs to use computer models [37] that are based on reasoning and follow the rules of thought.
The "Turing Test" (i.e. acting properly) stream says that when AI talks to people, it should act smart. For AI to do these things, it needs to be able to do human jobs at least as well as people [38]. The Turing Test [39] can be used to check these criteria.
Finally, the "Rational Agent" stream sees AI as a smart [40] or reasonable [11] agent1. This agent not only acts on its own, but also with the goal of getting to the most logically desirable result.

Another way to describe AI is to talk about intelligence in general and then use what you learn to make tools that are intelligent. Legg and Hutter [41] describe a way to measure intelligence by using psychology concepts, ideas of human intelligence, and intelligence tests. They use an agent-environment approach to talk about intelligence in general and artificial intelligence in particular when the agent is a machine. The "acting rationally" stream and their structure are very similar.

Aside from describing AI in general, another area of AI study is how to group AI into different groups. Searle [42] suggests telling the difference between weak AI and strong AI. A weak AI only acts like it can think, but a strong AI is a thought with feelings. However, Gubrud [43] sorts AI into groups based on the type of task. An artificial general intelligence (AGI) is an AI that can act at least as well as a human brain in any situation, without needing to be programmed to do so. being aware. A narrow AI, on the other hand, is one that can do as well as or better than the human brain in very particular jobs [44].

We will talk about the "Rational Agent" stream in more depth below because it is important for figuring out how to use machine learning in AI. The other three lines of study will be looked at again in Section 3, where we show that they work with our agent-based AI system.

“Rational Agent” stream says that intelligence shows up in the form of agents doing things. Five things describe these agents: they "operate autonomously, perceive their environment, persist over a long period of time, adapt to change, and set and pursue goals" [11, p. 4]. An agent decides what to do not for itself but with the people and things it deals with. Its sensors pick up on what's going on around it, an agent program figures out what to do with the information it gets, and its motors make it do something. To be a reasonable agent, the agent must also act in a way that leads to the best possible result based on this performance measure and what it knows about the environment and the actions that can be taken now and in the past.

As a general way to divide agents, Russell and Norvig say that the agent program can be split into four types of agents [11]: A model-based reflex agent also takes into account the agent's internal state, while a simple reflex agent only responds based on its sensing data. For an agent with goals, the best choice is the one that helps it reach those goals. The achievement of a goal is a simple choice, which means it can be done or not. A utility-based agent, on the other hand, doesn't have a clear goal; instead, it tries to maximise a whole utility function. By adding to its program, an agent can become a learning agent. Then, this kind of learning agent has a performance element that uses input data to decide what to do and a learning element that learns from its surroundings, creates its own problems, and tries to make the performance element better whenever it can.

There are three parts to the agent-environment framework: an agent, a setting, and a goal. This term refers to a "agent's ability to achieve goals in a wide range of environments" [41, p. 12]. The agent gets information from what it sees and hears in its surroundings. Observations of the world are one type of awareness. Reward signs that show how well the agent's goals are met are another type. The robot chooses what to do based on these signals, and the results are sent back to the world as signals.

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

When trying to understand the role of machine learning within AI, we need to take a perspective which has a focus on the implementation of intelligent agents. We require this perspective, as it allows us to map the different tasks and components of machine learning to the capabilities of intelligent agents. If we regard the capabilities of *thinking* and *acting* of an intelligent agent and translate this into the terms of software design, we can reason that the *acting* capabilities can be regarded as a *frontend*, while the *thinking* part can be regarded as a *backend*. Software engineers typically strictly separate form and function to allow for more flexibility and independence as well as to enable parallel development [45]. The frontend is the interface the environment interacts with. It can take many forms. In the case of intelligent agents it can be a very abstract, machine-readable web interface [46], a human-readable application [47] or even a humanoid template with elaborated expression capabilities [48]. For the frontend to interact with the environment, it requires two technical components; sensors and actuators. *Sensors* detect events or changes in the environment and forward the information via the frontend to the backend. For instance, they can read the temperature within an industrial production machine [49] or read visuals of an interaction with a human [50]. Actuators on the other hand are components that are responsible for moving and controlling a mechanism. While sensors just process information, actuators *act*, for instance by automatically buying stocks [51] or changing the facial expressions of a humanoid [52]. One could argue that the Turing test [39] takes place at the interaction of the environment with the frontend, more precisely the combination of sensors and actuators if one wants to test the agent's AI of *acting humanly*. Despite every frontend having sensors and actuators, it is not of importance for our work what the precise frontend looks like; it is only relevant to note that a backend-independent, encapsulated frontend exists.

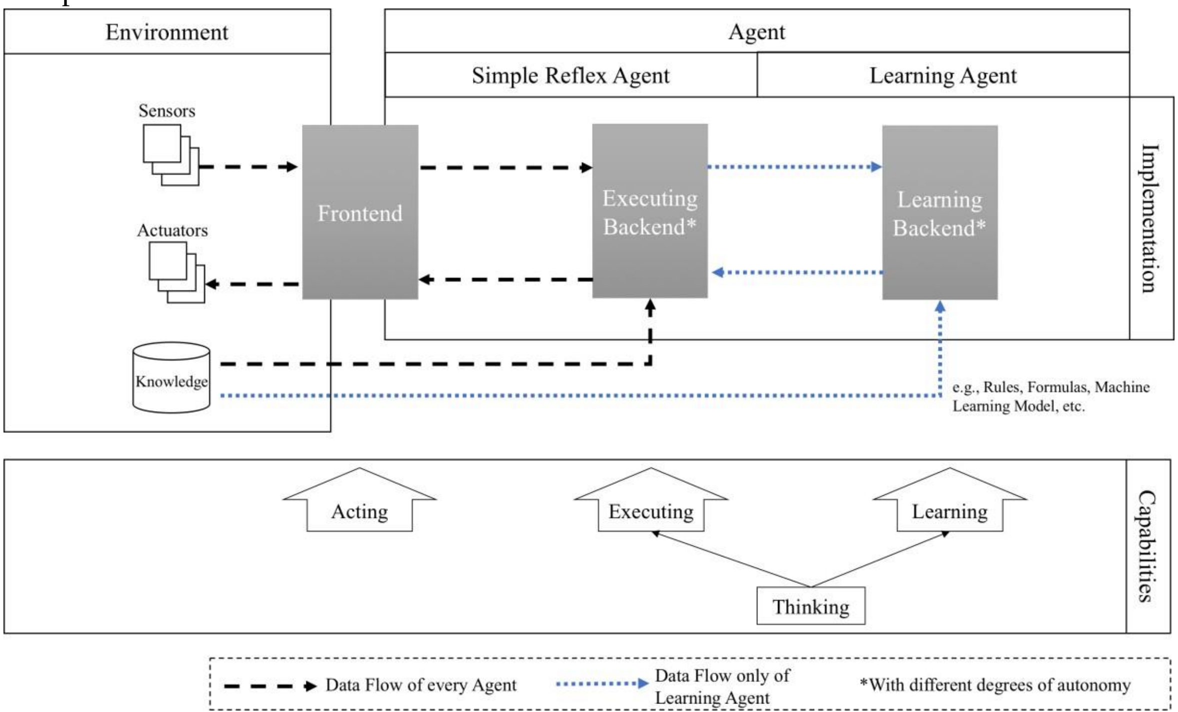

Figure 2. Conceptual framework

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

The learning backend dictates first *if* the intelligent agent is able to learn, and, second, *how* the agent is able to learn, e.g., which precise algorithms it uses, what type of data processing is applied, how concept drift [54] is handled, etc. Therefore, we pick up on the terminology from Russel & Norvig [11] by regarding two different types of intelligent agents: *simple-reflex agents* as well as *learning agents*. This differentiation especially holds for a machine learning perspective on AI, as it considers whether the underlying models in the *thinking* layer are once trained and never touched again (simple-reflex)—or continuously updated and adaptive (learning). In recent literature, suitable

examples for both can be found. As an example for simple-reflex agents, Oroszi and Ruhland build and deploy an early warning system of pneumonia in hospitals [55]: While building and testing the model for the agent shows convincing results, the adaptive learning of the system after deployment might be critical. Other examples of agents with single-trained models are common in different areas, for instance for anaphora resolutions [56], prediction of pedestrians [57] or object annotation [58]. On the other hand, recent literature also gives examples for learning agents. Mitchellet al. present the concept of "never- ending learning" agents [59] which have a strong focus on continuously building and updating models within agents. An example for such an agent is shown by Liebman et al., who build a self-learning agent for music playlist recommendations [60]. Other cases are for instance the regulation of heat pump thermostats [61], an agent to acquire collective knowledge over different tasks [62] or learning word meanings [63].
The choice on this feature in general (simple-reflex vs. learning agent) influences the overall design of the agent as well as the contribution of machine learning. The overview of our resulting framework is depicted in figure 2. In conclusion, in the case of a simple-reflex agent, machine learning takes places as a once-trained model in the execution sublayer. In contrast, it plays a role in the learning sublayer of a learning agent to continuously improve the model in the execution sublayer. This improvement is based on knowledge and feedback, which is derived from the environment via the execution layer.

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].

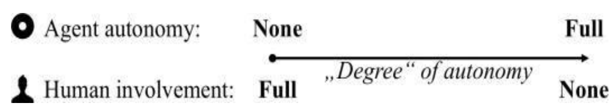

Figure 3. Degree of agent autonomy and human involvement

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

The presented framework of machine learning and its role within intelligent agents is still on a conceptual level. However, given the misunderstandings and ambiguity of the two terms [6–9], we see potential for further research with the aim both to clarify the terminology and to map uncharted territory for machine-learning enabled artificial intelligence.

First, empiric validation as well as continuous, iterative development of the framework is necessary. We need to identify various cases of intelligent agents across different disciplines and to evaluate how well the framework fits. It would be interesting to see how practical and academic machine-learning-enabled artificial intelligence projects map to the framework, and, furthermore even quantify which share of such projects works with learning agents and which with non-learning agents. Additionally, such cases would help us to gain a better understanding of the necessary human involvement in state-of-the art intelligent agents—and, therefore, determine the "degree" of autonomy when regarding all aspects (acting, executing, learning) of such agents.

Second, one aspect of interest would be to reduce the necessary involvement of humans. As stated before, we see this spectrum as a continuum between human involvement and agent autonomy. Two possibilities come immediately to mind. The methods of *transfer machine learning* deal with possibilities on how to transfer knowledge (i.e., models) from one source environment to a target environment [72]. This could indeed help to minimize human involvement, as further research in this field could show possibilities and application-oriented techniques to utilize transfer machine learning for automated adaption of novel or modified tasks [73].

Additionally, regarding already deployed models as part of the backend-layer, it is of interest not only how the models are built initially, but how to deal with changes in the environment. The so-called subfield of *concept drift* holds many possibilities on how to detect changes and adapt models—however, fields of successful application remain rare [54], [74].

5. Conclusion

In this study, we explain what machine learning does in the field of artificial intelligence, more specifically in the field of intelligent robots. We show a system that shows the two types of agents—simple-reflex and learning agents—and what machine learning can do in each of them. To sum up, machine learning models can be used in intelligent agents as models that have already been taught and can't learn anything new from their surroundings (simple reaction agent). When it comes to implementation, this part of the executing information is called the executing core. In this case, the robot can use machine learning models that have already been made, but it can't make and change its own. It is a learning agent, though, if the agent can learn from its surroundings and can then change the machine learning models in the processing sublayer. The learning core is an extra layer that learning agents have that lets them use machine learning for model building and training.

It is important to know how autonomous the machine learning in the agent needs to be when these two sublayers are being put in place. This part is all about how people need to be involved in machine learning jobs like gathering data or picking an algorithm.

The study being done right now is just an idea at this point, so it has some limits. First, the suggested framework can help us learn more about machine learning in AI, but more research is still needed to see how well current AI apps that use machine learning fit into this plan. Interviews with AI designers could be used to confirm the model and check the amount of detail. Also, we need to find ways to measure how much human input there is in AI jobs that involve machine learning so that we can get a better idea of how autonomous state-of-the-art bots really are. Even though it's still early, our approach should help scientists and professionals be more clear when they talk about AI and machine learning. It shows how important it is not to use the terms using them both in the same sentence, but making it clear what role machine learning plays in a certain agent application.

References

[1] "Transcript of Mark Zuckerberg's Senate hearing," The Washington Post, 2018. [Online]. Find the text of Mark Zuckerberg's Senate hearing at https://www.washingtonpost.com/news/theswitch/wp/2018/04/10/transcript-of-mark-zuckerbergs-senate-hearing/?utm_term=.4720e7f10b41.[Viewed on June 15, 2018].

[2] Z. Waseem and D. Hovy, "Hateful Symbols or Hateful People? How to Spot Hate Speech on Twitter: Predictive Features, in Proceedings of the 2016 NAACL Student Research Workshop, pp. 88–93.

[3] "Detecting Hate Speech on the World Wide Web," work by W. Warner and J. Hirschberg, presented at the LSM '12 Proc. Second Work. Language and Social Media, no. LSM, 2012. pp. 19–26.

[4] H. Chen, W. Chung, J. Qin, E. Reid, M. Sageman, and G. Weimann, "Uncovering the dark web: a case study of jihad on the web," International Journal of Research in Open and Distance Learning, vol. 14, no.4, 2013, pp. 90–103.

[5] H. Fujii and S. Managi, "Trends and priority shifts in artificial intelligence technology invention: A global patent analysis," Economic Journal, vol. 58,no.5, Pages 60–69, 2018.

[6] 2016 interview transcript from the University of Wisconsin with Mark Zuckerberg called "How to Build the Future."

[7] The Information Commissioner's Office spoke about "Big Data, AI, machine learning, and data protection" in their 2017 report on the Data security Act and the General Data Protection Regulation.[Online].This paper can be found at https://ico.org.uk/media/for-organisations/documents/2013559/bigdata-ai-ml-and-data-protection.pdf. [Accessed: June 15, 2018].

[8] "Big Data Management, Access, and Protection," Journal of the American College of Radiology, vol. 14, no. 5, pp. 579–580, 2017, by J. A. Brink.

[9] T. Nawrocki, P. D. Maldjian, S. E. Slasky, and S. An article by G. Contractor called "Artificial Intelligence and Radiology: Have Rumours of the Radiologist's Demise Been Greatly Exaggerated?" was published in Academic Radiology in 2018.

[10] C. F. Camerer, "Artificial intelligence and behavioural economics," in Economics of Artificial Intelligence, University of Chicago Press, 2017. Artificial Intelligence: A Modern Approach, 3rd ed. 2015, by S. J. Russell and P. Norvig.

[11] "The Elements of Statistical Learning: Data Mining, Inference, and Prediction," Math. Intell., vol. 27, no. 2, pp. 83–85, 2005, was written by T. Hastie, R. Tibshirani, J. Friedman, and J. Franklin.

[12] O. Bousquet, U. von Luxburg, and G. Rätsch, Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2–14, 2003, Tübingen,Germany, August revised lectures, vol. 3176, pages 4–16, 2003. Springer, 2011.

[13] Mohri, Rostamizadeh, and Talwalkar, The basics of machine learning. MIT Press, 2012.
[14] G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew in 2004 called "Extreme Learning Machine: a New Learning Scheme of Feedforward Neural Networks." What happened. 2004 IEEE International Joint Meeting Conference on, 2004, vol. 2, pp. 985–990.
[15] F. Sebastiani, "Machine learning in automated text categorisation," ACM Computer Society Journal, vol. 34, no. 1, 2002, pp. 1–47.
[16] LeCun, Y.,Bengio, Y., and Hinton,G. E.,"Deep learning," Nature,2015.
[17] 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 and was published in IEEE Signal Process. Mag. in 2012.
[18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
[19] C. Schommer, "An Unified Definition of Data Mining," CoRR, 2008, vol. abs/0809.2696.
[20] I. H. Witten, E. Frank, and M. a. Hall, Data Mining: Practical Machine Learning Tools and Techniques, Third Edition, vol. 54, no.2, 2011.
[21] Question marked as True: "What is the difference between AI, data mining, statistics, and machine learning?" 2014.
[22] J. R. Koza, F. H. Bennett, D. Andre, and M. A. Keane in Artificial Intelligence in Design '96 called "Automated Design of Both the Topology and Sizing of Analogue Electrical Circuits Using Genetic Programming."
[23] "Machine learning: Trends, perspectives, and prospects," By M. I. Jordan and T. M. Mitchell in Science in 2015.
[24] In Machine Learning,no.1,1997,T.M.Mitchell writes about this is from The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Vol. 9 by T. Hastie, R. Tibshirani and J.Friedman,published by Springer in 2017.
[25] R. Hirt, N. Kühl, and G. Satzger, "An end-to-end process model for supervised machine learning classification: from problem to deployment in information systems," in Proceedings of the DESRIST 2017 Research-in-Progress, 2017.
[26] Psychological Thinking by U.Neisser, 1967.
[27] N. J. Nilsson, Artificial Intelligence: A New Synthesis, vol. 125, no. 1–2, 1998.
[28] “The Logic of Action,” by K. Segerberg, J.-J. Meyer, and M. Kracht, in The Stanford Encyclopaedia of Philosophy (2016). [Online]. You can find it at https://plato.stanford.edu/archives/win2016/entries [Accessed: June 15, 2018]. /logic-action/.
[29] The 2016 Stanford Encyclopaedia of Philosophy has an article by R. Thomason called "Logic and Artificial Intelligence."
[30] "Getting ready for the future of AI" by the National Science and Technology Council and Penny Hill Press, Committee on Technology Vol. 58 from the National Science and Technology Council and Penny Hill Press. CreateSpace: An Independent Platform for Publishing, 2016.
[33] 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, J. Shah, "Artificial Intelligence and Life in 2030," by M. Tambe and A. Teller, was published in the Hundred Year Study Artificial Intelligence Report 2015-2016.
[34] IN 1961, A. Newell and H. A. Simon wrote "GPS, a program that simulates human thought."

[35] D. McDermott and E. Charniak, "Introduction to artificial intelligence," International Journal of Adaptive Control Signal Process, vol. 2, no.2, 1985,pp. 148–149.
[36] E.Rich and K. Knight, "Artificial intelligence," New from McGraw-Hill in1991.
[37] A. M. Turing, "Computing Machine and Intelligence," MIND, vol. LIX, no. 236, 1950, pp. 433–460.
[38] "Computational Intelligence and Knowledge," Comput. Intell. A Log. Approach, no. Ci, pp. 1–22, 1998, is a paper by D. L. Poole, A. Mackworth, and R. G. Goebel.